FearNot! - An Experiment in Emergent Narrative

R.S. Aylett[1], S. Louchart[2], J. Dias[3], A. Paiva[3], and M.Vala[3]

[1] MACS, Heriot-Watt University, Riccarton, Edinburgh EH10 4AS
ruth@macs.hw.ac.uk
[2] Centre for Virtual Environments, University of Salford, Salford M5 4WT
s.louchart@salford,ac.uk
[3] Instituto Superior Tecnico and INESC-ID, Av. Prof. Cavaco Silva, IST, Taguspark, Porto Salvo, Portugal
ana.paiva@inesc-id.pt, joao.dias@gaips.inesc-id.pt
marco.vala@tagus.ist.utl.pt

Abstract. We discuss the experience of constructing the application FearNot! (Fun with Empathic Agents Reaching Novel Outcomes in Teaching), an application of virtual drama to anti-bullying education inspired by Forum Theatre. An appraisal-driven agent architecture is presented as a mechanism for generating an emergent, that is, unscripted, narrative. A small-scale evaluation is discussed and the lessons learned are described.

1 Introduction

Virtual Storytelling (VS) has recently become an active research field in AI with enthusiastic researchers, active working groups and a growing community [6,12, 13, 15, 16, 18]. Although the VS community is now well established there are still many fundamental differences between approaches and frameworks and no generally agreed theoretical framework has as yet been established.

The concept of *Emergent narrative* [1, 2, 3,11] addresses the narrative paradox [3] observed in graphically represented VS. It revolves around the conflict between pre-authored narrative structures – especially plot - and the freedom a VE offers a user in physical movement and interaction, integral to a feeling of physical presence and immersion. The overall project could be described as the creation of a graphical system involving participating users in a highly flexible real-time environment where authorial activities are minimised and the distinction between authoring–time and presentation-time is substantially removed. Authorial activities would be limited to the set up of the story – in particular to the creation of characters and their milieu - and to providing the users with the necessary background information needed for them to play a significant part in the unfolding of the story. There would be no pre-determined end to the story or event time line, the development of the story would be handed to both the user and the Intelligent Agents and depend entirely on the interactions between themselves and their environment. The role of the author would thus be limited to the one of elaborating a high-level plot: this would be necessarily hypothetical in nature, since though it would be possible to have an idea of what the different characters could do, there would be no certainty that they would behave as expected.

T. Panayiotopoulos et al. (Eds.): IVA 2005, LNAI 3661, pp. 305 – 316, 2005.

The richness of characters and world together with advanced user interaction modalities needed to make this a reality are formidable, and as an initial step towards the overall concept, we describe here a much smaller-scale application of the ideas of emergent narrative in the demonstrator FearNot! – created as part of the EU-funded project VICTEC – Virtual ICT with Empathic Characters.

2 FearNot!

VICTEC, involving five partners in the UK, Germany and Portugal, sought to apply virtual dramas acted out by 3D graphically-embodied characters to what is known generically in the UK as Personal and Social Education (PSE) (or more recently as Personal, Social and Health Education – PSHE). This covers topics such as education against bullying and racism, on drugs, including smoking and alcohol, and sex education. A common thread in these topics is that knowledge in and of itself is not sufficient to meet the pedagogical objectives, since attitudes and emotions are at least as important to producing desired rather than undesired behaviour. For this reason, techniques such as small-group discussion, role-play and dramatic performance by Theatre-in-Education TiE) groups may be used.

A motivation for the project was to try to create some of the impact of dramatic performance through virtual dramas. The specific topic selected was anti-bullying education. Effective though theatrical performance is in this domain, it is necessarily collective, and in any group it is very likely that some individuals will be victims of bullying by some other in the group and thus will be inhibited in their participation. Thus a virtual drama application that could be used by the individual seemed to have a possible use.

The aim of the FearNot! (Fun with Empathic Agents Reaching Novel outcomes in Teaching) demonstrator was to allow children to explore what happens in bullying in an unthreatening environment in which they took responsibility for what happened to a victim, without themselves feeling victimized. The creation of an empathic relationship between child and character was seen as the mechanism through which this sense of responsibility would be achieved, so that the child user would really care what happened to the victimized character. The child was asked to act as an 'invisible friend', and to give advice which would influence the behaviour of the victim without undermining its autonomy of action and the child's ability to believe in it as a character with an independent inner life.

The interactional structure of FearNot! was inspired by the Forum Theatre approach developed by Brazilian dramatist Augusto Boal [5] in order to incorporate theatre into the development of political activism. In this dramatic form, an audience is split into groups, with each group taking responsibility for one of the characters in the drama. Between episodes of dramatic enactment, each group meets the actor, who stays in role, and negotiates with them what they should do next in the drama, respecting the constraints of their role and character. This structure of dramatic episodes divided by periods in which advice can be given to a character has been adopted for FearNot! as shown schematically in Figure 1.

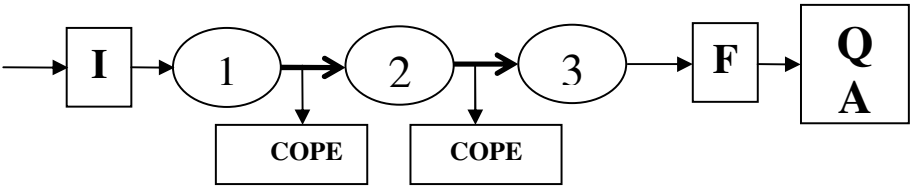

Fig. 1. Interactional structure of FearNot!

The session starts with an introduction to the school and the characters (I) and then a dramatic episode follows (1) in which a bullying incident occurs (see Figure 2 for an example). The victim then asks the child for advice in dealing with this, and the child suggests a coping behaviour (COPE). This structure is repeated – currently twice – and a simple educational message (F) is displayed, followed by an online questionnaire (QA) assessing how far the child can out itself in the shoes of the characters he or she has just seen.

The exploratory nature of the application is due to the lack of any 'magic wand' solution to the problem of bullying. Even the generally agreed educational message "Don't suffer in silence, tell someone you trust" is not guaranteed to work, though making a new friend and telling them is one of the more successful strategies. Some advice is controversial – parents often tell children to 'hit back' when faced by physical bullying, while teachers are universally opposed to violent responses. In fact 'hitting back' is statistically not often successful, but since it is memorable when it does succeed it is quite possibly over-reported [17].

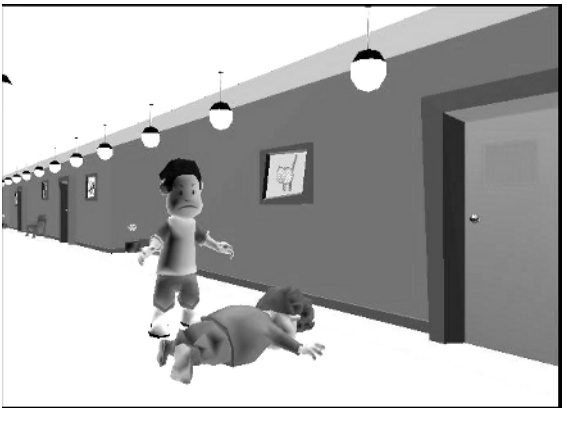

Fig. 2. A bullying incident

To retain the empathic link between child and victim, it is clearly helpful if the child feels the victim is taking the advice seriously. This is incompatible with a scripted approach, and indeed the use of a scripted version of the application in a large evaluation in June 2004, while demonstrating that children did indeed empathise with the characters, raised the criticism that the victim was not responsive to the advice given [9]. In this early version of FearNot!, only the third episode was influenced

by either of the two pieces of advice given. If one of these was to 'tell someone' then the victim was shown as improving their situation in the final episode, and if not, the third episode showed the situation was as bad as ever.

3 Narrative Management in FearNot!

Given there are around 7 different pieces of coping advice a child could give, and the order in which they are given before the second or third episode would also have to be taken into account, a branching narrative of the type used successfully in MRE [7] or Carmen's Bright IDEAS [12] seems infeasible. Thus an emergent narrative approach, in which action is driven by the characters themselves, is a natural solution to making the victim responsive to the advice the child gives.

At the same time, the repetitive nature of bullying, and the fact that it is naturally episodic, does not require too much from the emergent mechanism in terms of dramatic complexity or length. The Forum Theatre approach taken also means that the emergent mechanism does not have to take user actions directly into account. There were several good reasons for putting the child in the role of spectator during each dramatic episode. We have already mentioned the need to offer distance for children who are being bullied in real life; in addition the fact that child users would be able to hurt the virtual characters without being physically hurt themselves would have created a real imbalance in roles. 'God-like' intervention is not feasible in the real world either, and in any case the educational aim was to promote reflection, not to create a 'bash the bully' game.

The choice of an emergent narrative mechanism did not however remove the need for a narrative manager. Unlike a Forum Theatre production in which the action is temporally contiguous, it was always envisaged that each episode would be free-standing and could be thought of as happening over an extended period of weeks. Thus a choice has to be made about where each new episode is located and which characters are involved in it, as well as any other initial conditions. For example, it was envisaged that if the advice was to 'tell a teacher' or 'tell a parent'., then this would happen off-stage to avoid the difficult issues involved in representing teachers and parents as (possibly less than perfect) story characters. The initializing of episodes also allows a pedagogical influence to be exerted in terms of the situations and characters to be considered, which could be used for example to tailor FearNot! to specific schools. In addition, there has to be some method of determining when an episode has finished once there is no script encoding this information.

For these reasons, a Stage Manager agent was included in FearNot!, [15] with a series of levels of control, from determining every character action (used in the scripted version, and also in both versions for the introduction segment), to the ability to intervene in one-off fashion (used to end episodes), to merely setting up scene and characters and merely monitoring what happens.

4 An Affectively-Driven Architecture for Characters

If what happens in an episode is to be driven directly by character interaction, then a key aspect of the system must be the agent architecture. With an emergent narrative

mechanism, it is the ability of characters to autonomously select actions – their action-selection mechanism – that determines the narrative. Figure 3 shows the affectively-driven agent architecture.

Each agent in the world (the character) perceives the environment, through a set of sensors (allowing the perception of events, objects, etc. in the world) and acts on the environment though its effectors, allowing different actions to be performed (for example, a bully may hit the victim and the victim may cry). Upon receiving a percept (for example, be the presence of another agent or an object, or even an action from another agent) the agent appraises its significance and triggers the appropriate emotions. Additionally, if a goal has become active, it will add a new intention to achieve the active goal.

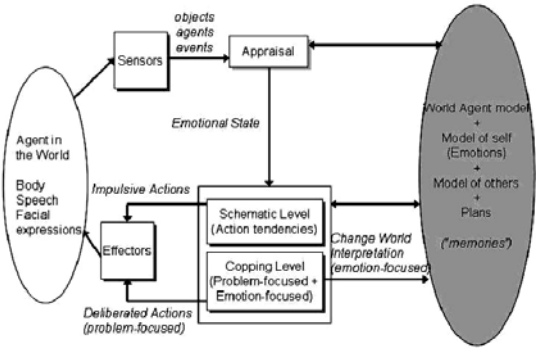

Fig. 3. Affectively-driven agent architecture

The appraisal process feeds the resulting emotional state into action-selection at two different levels: that of action-tendencies and that of coping behaviour [10]. For example, if the victim character starts to cry when bullied, it is not because s/he has a goal that involves crying – this is an innate reaction to a particular distressed emotional state and the inability to fight back.

On the other hand, other actions, such as begging the bully to stop, do result from the internal goals of the agent and are planned.

This second layer defines two kinds of coping: problem-focused coping, involving planning and acting to achieve goals; and emotion-focused coping in which the character's interpretation of the environment is altered. For example, an agent that feels distressed by being unable to achieve a given goal, may lower the goal's importance as a way of reducing its distress. In this way, emotions will not only influence the agents' reactive behavior, but also guide the planning process, since emotional focused coping changes the agent's interpretation of its plans. The continuous partially-ordered planner used in FearNot! selects the intention associated with the currently most intense emotion from the intention structure. This becomes the target goal for the planner to achieve. The planner then either removes a flaw or executes an action. The resulting plan is stored with the intention and can be pursued later.

5 Appraisal

The emotional component of the architecture applies a subset of the appraisal rules from the taxonomic definitions of Ortony, Clore and Collins (OCC) [14] which can be seen as a subjective evaluation of a given stimulus according to the character's goals, standards and beliefs. Figure 4 shows three examples of such rules, which generate the majority of OCC emotion types: *Well Being* emotions, *Attraction* emotions, *Fortune of Others* emotions and *Attribution* emotions, and are similar to those of [19].

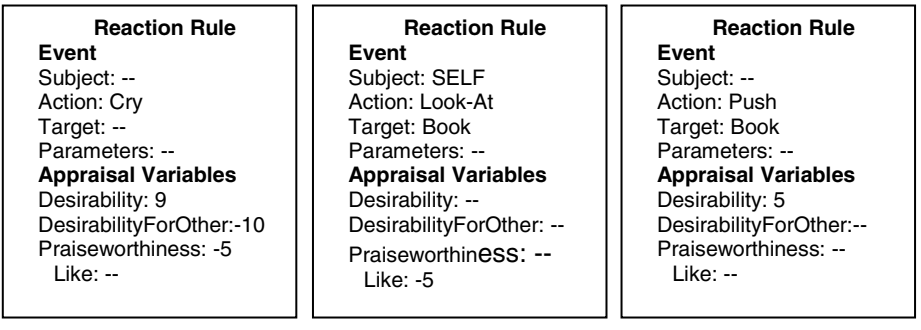

Fig. 4. Three examples of emotion appraisal rules

Table 1. Types of links between goals

Goal link	Description
SufficientTo	If goal A has a sufficient link to goal B then achieving A will also achieve B.
NecessaryTo	If goal A has a necessary link to goal B, then, in order to achieve B, one must achieve A
FacilitativeTo	If goal A has a facilitative link to goal B with value c, achieving A will raise the likelihood of achieving B by a factor of c.
InhibitoryTo	If goal A has a inhibitory link to goal B with value c, achieving A will lower the likelihood of achieving B by a factor of c.

Two of the OCC-defined goal types - *active-pursuit* goals and *interest* goals – are used. *Active-pursuit* goals are those the characters plan to achieve directly, such as physically attacking a victim. *Interest* goals are those a character has but does not actively pursue, such as avoiding getting hurt. Unlike the *active-pursuit* goal, the *inter-*

est goal does not have any pre-conditions, success or failure conditions since it does not become active or inactive. Instead it has a protection-constraint parameter, modeling those conditions that the character wishes to maintain. To allow the system to build a goal hierarchy, both goal types may possess several goal links as seen in Table 1.

The prospect-based emotions hope and fear are not however dealt with through domain-specific rules (as in [19], Hope and fear are related to goal achievement or not, so a similar approach to [8] was taken which takes advantage of explicitly storing the agent plan state and intentions. Prospect based reactions can then be automatically obtained from the plans and goals active in the agent memory.

6 Creating a Story

In this section we examine an example of an emergent narrative in order to show how the components already discussed fit together.

In the first episode, the Stage Manager locates John, the victim in the classroom studying and has Luke enter. Luke does not like John and so when he sees John he starts insulting him (reactive action tendency). As a result, John has an active pursuit goal of fighting back that is triggered when he is insulted by other characters. He tries to build a plan in order to fight back. However all the actions that John considers have some likelihood of getting hit back. When such an action is selected, a threat to John's interest goal of not getting hurt is detected and John feels frightened. Because he has a fearful nature (part of the personality profile for a victim), his fear is much stronger than the hope of succeeding in fighting back and so he gives up the goal and does not do anything.

Fig. 5. In the classroom

Fig. 6. User interaction

At the same time, Luke notices the book on the table and generates a bullying opportunity. He makes a plan to push John's books to the floor. Figure 5 shows a snapshot of this situation. Luke feels confident of his plan, so he starts walking towards the book with a happy face (the hope emotion is mapped to a happy facial expression). On the other hand John feels very distressed at being insulted and disappointed by not being able to fight back. Luke moves towards the books and pushes them

away. This event matches an emotional reaction generating the emotion *gloat*, which triggers an action tendency. Luke performs a *tease* language action that corresponds to saying something like: "Come and get them you Muppet!" When the victim realizes that the books are on the floor he activates the goal of picking them, and thus walks towards them and picks them up. When the bully sees John picking up the books he decides to push him. Once more this is achieved by an active pursuit goal that becomes active in that situation. So Luke goes behind John and pushes him.

The result of pushing John is uncertain: in the real world it is decided by physics, and in the virtual world by a probability set in the 3D visualization. Thus sometimes a character may fall, and in others, not. If John falls, he appraises this event as very undesirable and activates an action tendency to start crying. At the same time, Luke appraises the same event as very desirable and starts gloating John by saying something in the lines of "What a wimp, I've hardly touched you". When John cries, Luke finds it very blameworthy and thus threatens him to stop crying and to not tell anyone. If John does not fall, Luke will not mock him. Instead, the victim may feel angry and asks Luke why is he always picking on him. Luke responds negatively to the question by insulting John even more. Figure 6 shows a snapshot of the interaction mode in which the child user talks with the character victim and advises him/her on what to do next. The user types whatever he wants in the lower text box on the right and by pressing the OK button the written utterance is sent to the agent. The agent receives the utterance and converts it to a language action using a template-based language system [2]. When the interaction mode is first displayed, John arrives in the library crying, but he realizes that the user has entered the set as for any ordinary character (in fact the agent victim does not distinguishes the user from other synthetic agents) and activates the goal of asking for help which makes him perform an *askforhelp* speech act. If the user then suggests fighting back, this has the effect of raising the importance of the goal, so that the next time John meets Luke the fear generated by the possibility of getting hurt is not strong enough to make him give up the goal. Thus user interaction changes the behaviour of the victim by indirect influence rather than because the victim does exactly what he is told. However if John tries pushing Luke and it does not succeed, then he will not accept a further suggestion to hit back since the experience of being hurt as a result again alters his emotional state, this time in the direction of greater fearfulness.

7 Small-Scale Evaluation

A small-scale evaluation was carried out with eleven children randomly chosen from the third and fourth grade in a Portuguese school. The physical bullying story just described was used and each child participated individually. After the initial introduction and the first episode, each child was asked to write anything in order to help the victim. The victim had already asked for help, but the children did not always realize that they could really write something. All the interactions with the victim were saved in log files with a unique code for each child. At the end of the trial/interaction each child completed the same agent questionnaire that had been used in the large-scale evaluation of the scripted version of FearNot! [9]. One additional question was introduced relating to the dialog between child and victim, This could not have been used

with the scripted version since dialog was handled by menu selection in that version. It asks the child if the victim understood the conversation (by giving appropriate responses to the child's inputs).

Also differently from the scripted version, the emergent version has no sound at all. This is a disadvantage as the episodes may not seem so engaging, making the understanding of the story more difficult. Moreover, the lack of sound in the character dialogs requires the children to read the utterances written on the screen, which is more difficult than simply hearing them. Some children had difficulties reading utterances and in a few cases, they took so long to read a line that it disappeared before it was all read. In those few situations the researchers briefly explained what had been said. In terms of empathy with the characters, very similar results were obtained as with the scripted version: children disliked the bully and felt sad for the victim. However noticeably better results were obtained for aspects relating to the responsiveness of the characters as seen in Table 2. The first two questions refer to the conversation and dialogue between the characters.

Table 2. Responses to questions about character responsiveness

	Scripted version	Emergent Version
Conversations: did the conversations seem real? (yes-1;no-5)	2.4	1.9
Were the conversations (interesting-1; boring-5)	2	1.64
Did the victim understand the conversation? (yes-1; no-5)		1.36
Did the victim follow the advice? (yes-1; no-5)	2.3	1.7
Did you help the victim? (helped a lot-1; no- 5)	1.8	1.27

Since the episodes displayed are physical bullying episodes which contain few dialogue lines and the dialogues in the emergent version are very similar to the scripted version, the different results can be explained by the influence of the interaction with the character. The conversation with the victim makes the children look at the characters as more believable. For instance when the victim accepts the fight back strategy, it seems more real to see him threatening the bully on the next episode than to behave as in the first episode.

8 Lessons Learned

The first lesson of the work reported here is that a substantial amount of effort is required to produce an essentially bottom-up system. Because interaction between characters is the driving force for the development of narrative, the whole agent architec-

ture and the surrounding framework allowing agents to interact with each other have to be completely in place before any real testing of the narrative produced can be carried out. This is very different from a top-down approach in which a subset of facilities can typically be made available early and then elaborated. In particular, if emergent narrative is to be presented graphically, the graphic visualization must support full agent autonomy, including movement in the environment and the execution of animations. Due to the way in which the graphical world had been designed in WildTangent, autonomous characters were able to walk through furniture rather than around it, and in the absence of a viable implementation for local sensing in the WildTangent 3D world, waypoints had to be defined to support very simple path-planning.

In addition, when the character is itself able to decide what action to carry out, the animation that represents it in the graphical world must be visually correct, and this requires the character to position itself so that this is true. For example, if a push animation is designed such that the victim is pushed from behind, then it will only look correct visually if the character carrying it out is indeed standing behind the victim. In order that the character can check this before executing the animation, it was necessary to design spatially-specific execution points for animations, and include the necessary motion planning for a character to move to the correct execution point.

A further issue in the graphical environment is how to deal with dramatic cinematography when the actions and movement of characters are being decided on the fly. Camera position and lighting effects can make a great deal of difference to the dramatic impact of a scene on the user, and the scripted version was noticeably more competent in those respects. Once characters have autonomy, then the intelligence embedded in camera and lighting agents has also to be increased.

Speech output raises particular problem too in an unscripted environment. The template-based language system developed for FearNot! seems perfectly capable of generating the range of utterances needed for inter-character dialogue, and also coped – rather better than had been feared, and in both English and Portuguese – with character-child dialogue. However, given the robotic nature of text-to-speech synthesis systems, it was decided at an early stage to stick to text output on the screen rather than destroy the believability of the characters. Recorded speech would have been suitably expressive, but the amount of recording needed for the generative language system was prohibitive. Good quality unit-selection based speech systems are commercially available, but they currently require the load into memory of a very large database – incompatible with the resources available when running interactive graphics – and moreover have been designed for adult voices only and the equable tones of the telephone help system, not the angry or miserable child characters of FearNot!

A methodological point was raised by the use of this approach in an educational application. To what extent is the necessarily somewhat unpredictable outcome of episodes in conflict with the pedagogical objectives? It is possible for example for the Stage Manager to bring characters together with a view to bullying taking place and for none to happen. This is like the real world, but an educational application is more constrained than the real world. The use of the Stage Manager allows the degree of emergent narrative to be constrained if desired, and it may be that the amount of narrative variability that is acceptable will depend on the exact application chosen.

9 Related Work and Conclusions

In a project covering as much ground as this one, many pieces of previous work had an influence. We have already cited earlier work using an OCC approach such as [19] , while the use of an emotion-driven planner and of expressive behaviour for dramatic purposes can be seen in [8] [9] and [12]. Like most other researchers in this field we must also acknowledge the seminal work of the OZ project and in particular its emphasis on believability [4]. However, apart from the novelty of the application domain – no previous autonomous agent application has targeted anti-bullying education – the emergent narrative experiment was also truly novel in our view. Much other interesting story-telling work is going on, but no other group seems as yet to have attempted an unscripted approach in this way. Variation in story outcome has been generated for example by [6] but this is derived from pre-built goal-trees which interact in different ways for an initial random positioning of characters in an environment rather than generatively as in this case. Façade [13] is a beautifully designed story environment, but its conception of beats is closer to that of universal plans and produces a very large authoring task that may not be sustainable for an educational environment. VICTEC and Façade differ in their narrative approaches, the stories in VICTEC being created form rather than articulated around the user actions, as it is the case in Façade. The Mimesis environment [18] is a very interesting application of planning, but is aimed at authoring and not at unscripted drama as is [16]. As we have discussed in the previous section, many issues have arisen from the emergent narrative work carried out in FearNot! and further research is required to deal with these. However we believe that we have shown there is an interesting role for this approach to unscripted narrative, and that there may be applications such as this in which an open-ended and somewhat unpredictable narrative has much to offer.

References

1. Aylett, R (2000) Emergent Narrative, Social Immersion and "Storification"; Proceedings Narrative andLearning Environments Conference NILE00 Edinburgh, Scotland.
2. Louchart S, Romano D, Aylett R, Pickering J (2004) Speaking and acting - Interacting language and action for an expressive character. Proceedings for the AISB workshop Leeds Uk 2004
3. Aylett R, Louchart S. (2003) Towards a narrative theory of VR. Virtual Reality Journal, special issue on storytelling. To appear 2003.
4. Bates, J. (1994). The Role of Emotion in Believable Agents *Communications of the ACM* 37, 7 122—125
5. Boal, Augusto. (1979) The Theatre of the Oppressed. Urizen Books, 1979
6. Cavazza, M, F. Charles, S.J. Mead, Character-Based Interactive Storytelling, IEEE Intelligent Systems, July/August 2002, pp 17-24
7. Gratch J., Marsella S., (2001). Tears and Fears: Modeling emotions and emotional behaviors in synthetic agents, Proc. of Autonomous Agents 2001, Montreal, ACM Press. 278-285. 2001.
8. Gratch.J (2000) Emile: Marshalling passions in training and education. In *Autonomous Agents 2000*. ACM Press, 2000.

9. Hall, L; Woods, S; Dautenhahn, K; Sobral, D; Paiva, A; Wolke, D. and Newall, L. (2004) Designing empathic agents: Adults vs Kids.In: James C. Lester, Rosa Maria Vicari, Fábio Paraguaçu (Eds.): *Intelligent Tutoring Systems, 7th International Conference, ITS 2004,* (604-613) LNCS 3220 Springer
10. Lazarus, R. (1991) *Emotion and Adaptation*. Oxford University Press, 1991.
11. Louchart S, & Aylett R.S (2002) Narrative theories and emergent interactive narrative; Proceedings NILE02 Edinburgh, Scotland pp 1-8
12. Marsella, S; Johnson, L.W. & LaBore, C. (2000) Interactive Pedagogical Drama Proceedings, Autonomous Agents 2000
13. Mateas, M., & Stern, A. (2000). Towards integrating plot and character for interactive drama. In Working notes of the Social Intelligent Agents: The Human in the Loop Symposium. AAAI Fall Symposium Series. Menlo Park, CA: AAAI Press. A shorter version of this paper appears in K. Dautenhahn (Ed.), Socially Intelligent Agents: The Human in the Loop. Kluwer. 2002.
14. Ortony, G. Clore, and A. Collins. (1988) *The Cognitive Structure of Emotions*. Cambridge University Press, New York, reprinted 1994 edition, 1988.
15. D. Sobral, I. Machado, and A. Paiva. Managing authorship in plot conduction. In *Virtual Storytelling: Using Virtual Reality Technologies for Storytelling*. Springer,2003.
16. Szilas.N, 2003. *Idtension, a narrative engine for Interactive drama*, Proceedings TIDSE. 2003.
17. D. Wolke, S. Woods, H. Schultz, and K. Stanford. Bullying and victimization of primary school children in soth england and south germany: Prevalence andschool factors. British Journal pf Psychology, 92, 2001.
18. Young, R.M. (2001) An Overview of the Mimesis Architecture: Integrating intelligent narrative Control into an existing gaming environment. Working notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment, AAAI Press 2001
19. Martinho C.: "*Emotions in Motion: short time development of believable pathematic agents in intelligent virtual environments*", Universidade Técnica de Lisboa, Instituto Superior Técnico, Lisboa, Master Thesis, 1999.