Unscripted Narrative for affectively driven characters

Ruth Aylett[2], Sandy Louchart[1], Joao Dias[3], Ana Paiva[3], Marco Vala[3], Sarah Woods[4], Lynne Hall[5]

(1)CVE, Univ. Salford
Manchester M5 4WT UK
S.Louchart@Salford.ac.uk

(2)MACS, Heriot-Watt University
Edinburgh E14 4AS UK
Ruth@Macs.hw.ac.uk

(3)INESC-ID
IST, Taguspark
Porto Salvo, Portugal
Ana.paiva@inesc-id.pt, joao.dias@gaips.inesc-id.pt, marco.vala@tagus.ist.utl.pt

(4)Adaptive Systems Group
STRI, Univ. Hertfordshire
College Lane, Hatfield
Herts, AL10 9AB, UK
s.n.woods@herts.ac.uk

(5) Schl Comp. & Tech
Univ. Sunderland, PO Box 299, Sunderland, SR6 0DD, UK
Lynne.hall@sunderland.ac.uk

Abstract

The paper presents requirements for the design of unscripted (emergent) dramas based on research into role-playing games. It considers the FearNot! demonstrator in anti-bullying education as a sample implementation, describing the architecture of its affectively driven intelligent autonomous characters. It presents a comparative evaluation of the unscripted version against an earlier scripted version, examines related work and further development of the emergent narrative approach.
Keywords – AI, Autonomous agents, emotion modelling

Introduction

The video games industry has successfully demonstrated over the last two decades that virtual characters, virtual worlds/environments and even virtual societies can reach and entertain very large numbers of end-users. However, with some exceptions (for example The Sims, Black and White) much of the development effort over this period has gone into improving the quality of the graphics rather than developing more functional or autonomous characters, or new approaches to interactive narrative. Direct action has been prioritized and the current representation of narrative in today's video games (a cinematic, tree-based approach) has become a means of invoking action sequences rather than relating to the story experience or any overall narrative drive. However, the emergence of educational applications and the developing interest of the education community in the use of Virtual Reality (VR) technologies has raised important issues of narrative articulation and story representation in virtual environments. Since action is less of a focus in educational applications, the main focus here has switched from the linking of direct action sequences involving the player/user to the smooth articulation of role-play and educative content (i.e. role-playing, storytelling and participative activities).

The research reported here on the Emergent Narrative concept [1] aims at the definition of a narrative theory adapted to the VR medium (whether game or VR application). The inherent freedom of movement proper to VR, an indisputable element of immersion, collides with the Aristotelian [2] vision of articulated plot events with respect to the given timeline associated with the story in display. This narrative paradox can only be observed in interactive VR applications and it does not seem possible to resolve it through the use of existing narrative theories. From Plato's story definition [1] and Aristotle's plot consideration, all the way to Propp's meta-structural narrative articulation [3], Campbell's cyclical diagrams [4] and Barthes and the French Post-Structuralism's top-down analytical masterpiece [5], one novel element must be confronted: interactivity.

A form in which the audience is not static and has the option of interacting with the characters or environments of the story brings another dimension to storytelling altogether and extends the boundaries of both narrative creation and articulation. A story can from now on can not only be told by the author/poet directly (Diegesis, in Plato's terminology [1]) or shown to the audience through the use of characters (Mimesis), but, also be experienced and lived through – and the user cannot be forced to be where a specific plot is happening and do what it requires without destroying their suspension of disbelief and sense of immersion in the virtual world.

The emergent narrative solution

The emergent narrative concept presented in this work is based on the idea that a story, as well as being authored and displayed in the classical way, can *under specific* *conditions* that are the subject of this investigation, also emerge directly from the interactions between its different protagonists and build itself from the causal relationships between its different elements.

Since narrative theories do not currently deal with interactivity and present a rather restrictive consideration of the user with respect to story articulation, we have studied forms that break with the plot-directed stance and are more participative and interactive in their narrative approach. This includes improvisational and interactive forms of theatre such as Boal's Forum Theatre and street theatre [6], Reality TV, the role of Non-Player Characters (NPCs) in video games and in particular Role Playing Games (RPGs) in their many forms and aspects (i.e. cooperating and conflicting board-based, and live).

Research into RPGs was mainly empirical since though there is some practitioner literature, there are very few scholarly resources available for their detailed study. Using the approach of knowledge engineering, three experts, one in cooperative board-based RPGs ('cooperative' refers to the party carrying out the RPG quest), one in conflicting board-based RPGs (where the party have conflicting goals) and one in live RPGs were interviewed in depth using a professional knowledge acquisition system: KAT Builder [19]. Expert1 has more than 20 years experience as gamesmaster in cooperative board-based RPGs and was a 2001 world champion in an international RPG competition; expert2 also has more than 20 years experience in running short competitive RPGs; expert 3 has more than 5 years experience as a full-time professional working for a local authority on educational live RPGs. The KA tool was used to extract story management rules, discussed in more detail in [12] – a subset can be seen in Table 1.

In an emergent narrative, narrative unfolding and its significance are integrated threads of a single process, made of narrative tensions, causal links, logical and affective decisions, personalities and priorities. Thus it replaces the artefact-based view of narrative with a process-based view. Most of these elements are inherent to the characters, users or not, and replace the concept of the single protagonist at the centre of the story. Thus a character-based approach is not composed of a single storyline to which the different characters must conform in order to give sense to it, but of as many storylines as there are characters. It is this multiplicity of storylines that makes it a suitable approach for interactive drama and interactive experiencing. Although such an approach and associated techniques function well in the world of RPGs it must be adapted to direct computational implementation.

From a theoretical point of view, the articulation of a process-based narrative model suggests a greater value for multiple character-based experiences over less scalable plot and tree-type approaches. From a more practical perspective, since each character is in the centre of his/her narrative, the focus of development in an actual implementation must be oriented towards the completeness of the character.

Emergent narrative does not abolish authoring: rather it changes its requirements. The author must create roles, environments, props and relationships according to a global vision of the whole experience rather than a linear plot. If characters are to interact intelligently and meaningfully between themselves, their different potential relationships with each other must have been thought through and their place in the world must have been clearly established. In addition to the creation of both worlds and characters that would be likely to interact in interesting and potentially dramatic ways, the author must formulate the setting up and emergence of situations likely to trigger the different protagonists into action and decision-making. This raises the question of dynamic narrative control within this approach. While – as in most expert system projects - it is not yet feasible to model the whole of a human Game Master's (GM) expertise, the knowledge acquisition exercise already mentioned showed that elements of this expertise, in the form of techniques used by a GM to control the game's unfolding, can be modelled.

Context	Trigger type	Event type
Character Management		
The player is not interacting / not attentive	No interaction when there is an opportunity	Send an NPC to directly interact with the player and prompt a reaction
The player is suicidal	Player is taking obvious and unnecessary risks	Remind character of potential consequences of its actions
Drama Management		
Action takes longer than expected	The player has insufficient information to proceed	Send NPC to assess knowledge and highlight gaps (hints)
Unexpected branching of the story	The player is acting out of role	Remind roles and rules, bring next encounter
Player incorrectly determines what to do next	Player pursues wrong goal, goes in the wrong direction	Give hints they are going the wrong way, or emptiness

Table 1: Examples of character and drama trigger and event types

While a rich definition of character and modelling of elements of GM narrative control are essential to the success of the research presented, the definition of an autonomous affective agent framework is the key to its realization and implementation. Since the character is at the centre of narrative development from both its own and the system's perspective, the development of intelligent agents that can react and therefore act autonomously under certain stimuli (narrative, emotional, personality-related) is a requirement for translating theory into implementation. Affect is seen as central to the creation of unscripted narrative since it both produces dramatically interesting action-selection and the accompanying expressive behaviour required to establish the context of an action in a character's motivations. In the rest of this paper we discuss an initial experiment in implementing emergent narrative, carried out in the European Union Framework V project VICTEC (Virtual ICT with Empathic Characters). The agent framework developed for the project allows the construction of virtual intelligent agents that express and react to emotions in a natural and meaningful way. It has been designed so that it does not only apply to the specific context of school bullying, but can be used in the more general realization of emergent dramas.

The FearNot! application

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). 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 in producing desired behaviour. For this reason, techniques such as small-group discussion, role-play and dramatic performance by Theatre-in-Education TiE) groups may be used.

The project aim was to create some of the impact of dramatic performance through virtual dramas. The specific topic selected was anti-bullying education. Effective though TiE 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 and will be inhibited in their participation. This suggested a virtual drama application that could be used by the individual.

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 autonomous 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.

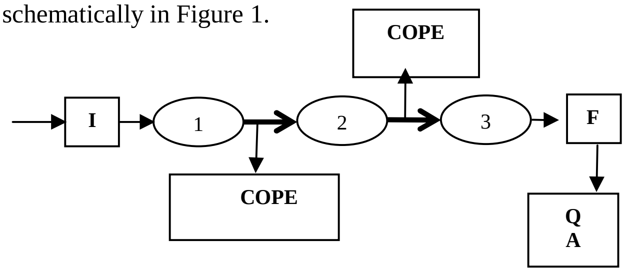

Figure 1: Interactional structure of FearNot!

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, but with a target of five or six episodes – and a simple educational message (F) is displayed, followed by an online questionnaire (QA) assessing how far the child can put itself in the shoes of the characters just seen.

The FearNot! agent framework

The agent architecture used in the FearNot! demonstrator is shown in **Figure 3.** Since the events during an episode are to be driven by character interaction, the appraisal-driven agent architecture forms a central part of the system. With an emergent narrative mechanism, it is the ability of characters to autonomously decide upon their own actions – their action-selection mechanism – that determines the narrative. 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

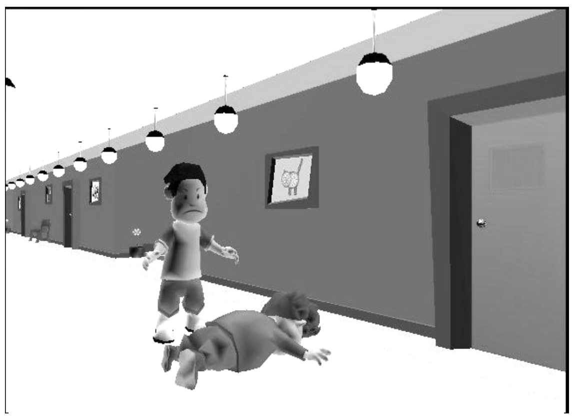

Figure 2: A bullying incident

victim and the victim may cry). Upon receiving a percept (for example, 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.

The agents' behaviours, rather than being generated by a conventional planner are fundamentally influenced by their emotional states and personality. Their emotional status affects their drives, motivations, priorities and relationships. FearNot! provides two distinct levels in both appraisal and coping mechanisms. The reactive level provides a fast mechanism to appraise and react to a given event, while the deliberative level takes longer to react but allows a much more complex and rich behaviour.

The appraisal process feeds the resulting emotional state into action-selection at two different levels: that of action-tendencies and that of coping behaviour. 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.

The emotion model

The emotion definition adopted is that of Ortony, Clore and Collins (OCC) [7]. The OCC model is an approach based on appraisal of the affective valence (good or bad) and intensity of the impact of an event, and the classification of emotions it defines can be seen as a hierarchical taxonomy of 22 emotion types. Using this model, emotions are represented in the architecture by the following attributes: **[Table 2]**

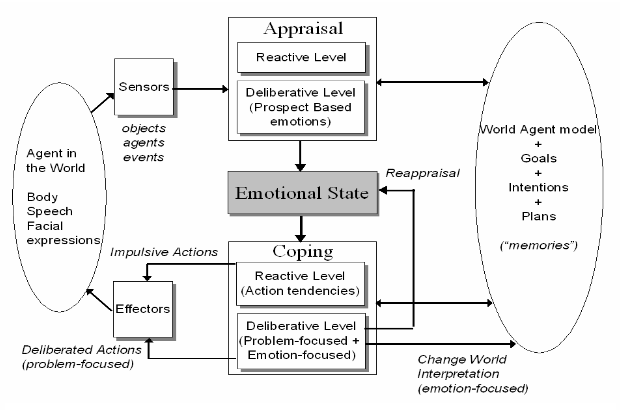

Figure 3: The FearNot! affective agent architecture

Each emotion type can be realized in a variety of related forms with varying degrees of intensity (i.e. emotion type Fear can generate an emotion range from concern to petrified). The attribute Valence describes the value, positive or negative, of the reaction that originated the emotion, while the target and cause attributes help in addressing and accessing both emotional impact and potential answer to the stimulus.

Attribute	Description
Type	The type of the emotion being experienced
Valence	Denotes the basic types of emotional response (positive or negative)
Target	The name of the agent/object towards the emotion is directed
Cause	The event/action that caused the emotion
Intensity	The intensity of the emotion. A logarithmic scale between 0-10
Time-stamp	Moment in time when the emotion was created

Table 2: VICTEC agents' emotion attribute

As a dynamic process, the intensity of an emotion must be attenuated through time from the moment it is generated onwards. This was modelled using Picard's [8] decay function for emotions, where intensity is characterized as a function of time. At any time (t), the value for the intensity of an emotion (em) is given by the formula:

$$Intensity(em,t) = Intensity(em, t0) \cdot e^{-d \cdot (t-t0)}$$

The value *d* (decay) determines how fast the intensity of this particular emotion will decrease over time. The value *Intensity(em, t0)*, refers to the value of the intensity parameter of the emotion (*em*) when it was created.

OCC can be used not only for a character's appraisal system, but also to define a character's personality. This is set in our model by specifying for each character, using XML:

- emotional thresholds and decay rates for each of the
- 22 emotion types defined in OCC.
- the character's goals;
- a set of emotional reaction rules;
- the character's action tendencies;

Ortony et al associate a threshold and decay rate with each emotion type, where the threshold specifies a character's resistance to an emotion type, and the decay rate, as before, specifies how fast the emotion decays over time.

According to OCC, when an event is appraised, the created emotions are not necessarily "felt" by the character. While the appraisal process determines the potential of emotions, such emotions are added to the character's emotional state only if their potential surpasses a defined threshold (specific to each emotion). Even where an emotion passes the threshold, the final intensity is given by the difference between the threshold and the initial potential:

$$Intensity(em) = Potential(em) - EmotionThreshold(em)$$

Thus a calm character will have a high threshold and a strong decay for the emotion type of *Anger*, will thus rarely experience anger, and then with lower intensities and for a short period of time. It is possible to have two characters with the same goals, standards and behaviours that react differently to the same event (by having different thresholds).

Operator
Action: GetUp
Preconditions:
  Property: ?SELF(Status) Op: != Value: Stand
Effects:
  Prob: 1.0 Property: ?SELF(Status) Value: Stand

Figure 4: Operator example

The appraisal mechanism

As shown in *Figure 1*, the appraisal mechanism is composed of two distinct layers. The reactive layer appraisal is handled by a set of emotional reaction rules, based on Elliot's Construal Theory [8]. A reaction rule consists of an event that triggers the rule and values for OCC appraisal variables affected by the event (desirability, desirability for other, praiseworthiness etc).

The deliberative layer is responsible for appraising events according to the character's goals, thus generating OCC prospect-based emotions like hope and fear. FearNot! includes two of the OCC goal types; *active-pursuit* and *interest goals*. *Active-pursuit* goals are goals that the character actively tries to achieve (i.e. going to a dental appointment) while *interest goals* represent goals that the character has but does not pursue (i.e. avoiding getting hurt). The OCC *replenishment goals* are not used since they could be considered as *active-pursuit goals* with cyclic activation and deactivation.

When an event is appraised, the deliberative level checks if any goal has become active, and if so, an intention to achieve the goal's success conditions is created, generating hope and fear emotions according to the goal's probability of success. At the same time, this layer monitors all active goals and actions chosen to achieve them, updating the probability of action effects thus changing plan probabilities and generating new hope/fear emotions.

The action selection and coping mechanism

Like the appraisal mechanism, the action selection process is composed of reactive and deliberative levels.

The reactive layer consists of a set of action rules: each contains a set of preconditions that must be true in order to execute the action and an eliciting emotion that triggers this particular action. The action set is matched against all the emotions present in the character's current emotional state and the set of rules with positive matches is activated. The action rule triggered by the most intense emotion is then selected for execution. If more than one action rule is selected (i.e. triggered by the same emotion), the most specific one is preferred.

To build the core of the deliberative or coping layer we implemented a standard partial-order continuous planner [14] and extended it to include emotion-focused coping strategies in addition to the common planning operations [22]. The planner selects the currently most intense intention, which corresponds to the goal generating the most intense fear or hope emotion after appraisal has been carried out.

More than one plan may be generated for this target goal and the planner selects one in order to continue planning or execution. Once the selected plan is brought into focus it generates hope/fear emotions, including emotions caused by action threats to interest goals. The continuous planner will then either remove a plan flaw or execute an action if the plan is complete.

Plans are modelled as a set of operators along with ordering constraints, causal links, binding constraints and open pre-conditions.. Operators are a slight modification of STRIPS [9] operators, associating probability values with the effects. **Figure 4** shows the operator for the *GetUp* action. In order to get up, the agent cannot be already standing up. This condition is represented by the character status property, which has to be different from "Stand". Since this property has one of the three values: "Stand", "LieDown" or "Sit", the character may get up if it is seated or lying down.

Emotion Focused Coping

Marsella and Gratch [18] introduced the use of emotionally focused coping in planning processes. This works by changing the agent's interpretation of circumstances, lowering strong negative emotions, and is often found in humans, especially when problem-focused coping (acting on the environment) has low chances of success.

The FearNot! deliberative planner uses the emotion focused strategies of *acceptance*, *denial* and *mental disengagement*. Acceptance is the recognition that something cannot be achieved, so that failure is accepted. When a plan has a very low probability of success, the planner will accept plan failure and will not try to improve it. If no other plan that achieves the goal remains, the goal also fails. But the most important role of acceptance is when a plan step threatens another goal (say an interest goal protected condition). If the active pursuit goal generates stronger emotions than the threatened interest goal, the plan is maintained and the protected condition failure is accepted. Otherwise, the plan will be dropped. Mental disengagement is used whenever acceptance is applied and works by lowering the goal's importance (thus lowering the disappointment experienced by the character).

Traditional planners deal with threats by applying promotion or demotion, i.e. by moving the threatening step to be before or after the threatened step. In addition to this process, the deliberative layer can use denial to deal with such threats. If the step effect that threatens the condition does not have a very high probability of happening, the agent can ignore the threat assuming that the effect will never happen by lowering the effect probability.

The agent architecture we have discussed here, where the planning and coping system are affectively driven, offers a useful test platform for the computational implementation of the emergent narrative concept described earlier. Indeed, since the agents are emotionally driven, any significant interaction with a child user or another agent will result in the alteration of the agent's emotional state.

Since the agent makes decisions based on that emotional state, this potentially affects its perception of actions and alters the probability of plan success and the resulting feelings of hope and fear. This, in turn, influences the actions selected for execution by the agent and allows for the unfolding of narratives different in form and content (i.e. according to their context) without the need for scripting them.

The FearNot! visualisation system

Figure 1 shows only a part of the FearNot! architecture: the 'agent mind' that controls the behaviour of a character and the virtual world model to which it is linked by *sensors* and *effectors*. Underlying this is a client-server architecture [**Figure 5**] in which control modules are linked to the world model via sensors and effectors handled by a generic message-passing system. In this generic framework, the agent minds, the user (just another agent), a Stage Manager (responsible for initialising episodes) and a View Manager all operate around the symbolically-represented central world model [**Figure 6**].

This virtual world that the characters inhabit is presented to the user through a visualization system. The View Manager, a special agent (which does not represent a character) with the power to listen to all the events that occur in the virtual world, has the task of translating those events to a specific visualisation system, which in the case of FearNot! was the game engine WildTangent. In this process, the characters' actions (expressed through effectors) are translated into a sequence of view actions. The parameters passed via the effector are used to specify the way the character will be visualized - for example the character's facial expression – as well as to map planner actions on to animations defined in the visualisation system. This modular separation of AI and graphics components makes each editable or replaceable: WildTangent could be replaced by, for example, Unreal Tournament, or even a 2D engine without having to reimplement the virtual world or agent minds. Only that part of the View agent oriented specifically to WildTangent would have to be changed.

This approach is very different from that commonly

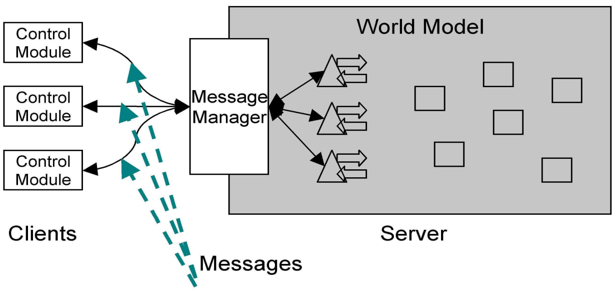

Figure 5: Client-server architecture

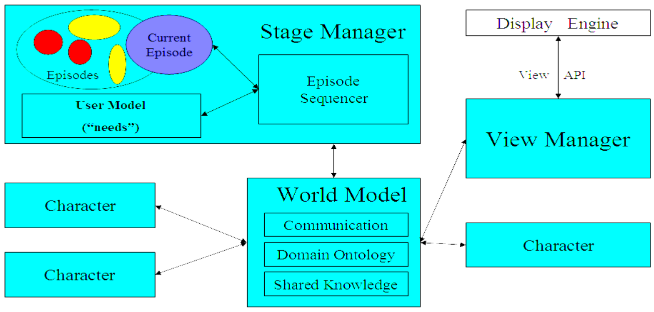

Figure 6: FearNot! overall architecture

employed by graphics researchers in which the whole architecture is embedded in the visualisation system, for example using Unreal Tournament scripting: it represents an extension and generalisation of work such as Gamebots [15] in which agent minds are directly attached to a visualisation engine via sockets. In our work an added virtual world model is the central and key component of the architecture and message –passing is generalised.

Of course the character bodies and animations, as well as the 3D graphic visualisation of the rest of the virtual world, have to be specified in a format the visualisation engine will accept and loaded in at run-time. Specific issues relating to 3D space also have to be dealt with in the visualisation engine, including path planning of the movements specified by the character minds. Additionally, in FearNot!, the outcome of actions which are not certain to succeed are also decided in the visualisation engine: for example whether a character who is pushed actually falls. The rationale for this design decision is that such outcomes depend on the physics of 3D space, not the logical relationships of the virtual world model.

Generating a story within FearNot!

In this section we examine an example of an emergent narrative – the scenario used in the evaluation discussed below - 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.

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 7** shows a snapshot of this

Figure 7: In the classroom

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 up, 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 visualisation - sometimes a character may fall, and sometimes, 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 by saying something like “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, John may feel angry and ask Luke why he is always picking on him. Luke responds negatively to the question by insulting John even more.

Figure 8 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 - one of the coping responses in **Figure 9** - using a template-based language system [11]. When the interaction mode is first displayed, John arrives in the library crying, but he realizes that the user has entered the set (the victim character does not distinguish the user from other synthetic agent) 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.

Evaluating FearNot!

A large scale evaluation (N: 401 children) of a scripted version of FearNot! was carried out at the University of Hertfordshire (2004), in a large IT suite for a period of two weeks, with up to 65 grade3-4 children participating each day [16]. The aim was to evaluate the overall pedagogical concept and the scenario material used before adding the autonomous agent architecture. Children were seated individually at work-stations and the project coordinator provided an overview of the day’s activities. During the session, children interacted with two bullying scenarios depicting physical and relational bullying incidents.

The interaction structure was as in Figure 1 above. In this version of FearNot!, the child was asked to select an advice from a randomly ordered list of coping strategies shown as a drop down menu **[Figure 9]** rather than entering free text as in the later unscripted version. This dialogue also asked the child to explain his/her selection and what he/she thought would happen after having implemented the selected strategy, by typing it in. At the end of the scenario, a universal educational message was displayed saying that “telling someone you trust” is usually a good choice. A significant difference between the scripted and emergent version of FearNot! was that in the scripted version, the action taken by the victim character to cope with the bullying situation, was not influenced by, or

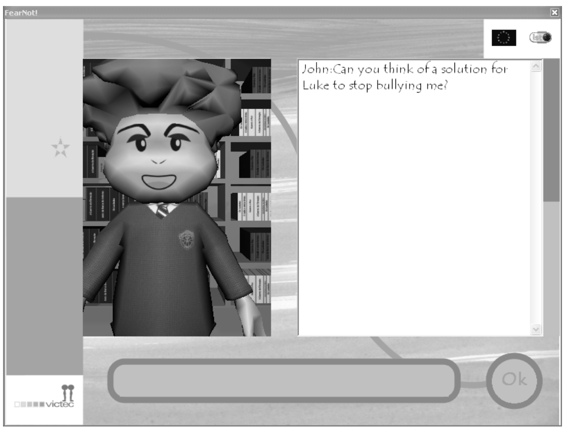

Figure 8: The child interacts with the victim

dependent on, the child's choice of coping strategy.

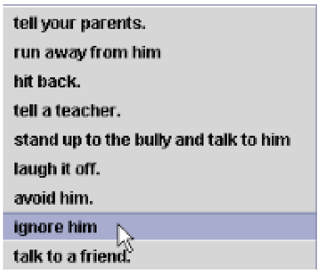

Figure 9: Coping choices

After children had interacted with FearNot!, they completed an Agent Evaluation Questionnaire (AEQ), and participated in a focus group activity. The structure of the questionnaire is shown in Table 4, and children's views were predominantly measured according to a 5 point Likert scale The AEQ was systematically developed throughout the course of the project and extreme care was taken to ensure that children understood the terminology and language used and the nature of the rating scale. Several pilot studies using the AEQ were carried out prior to the large scale evaluation event).

Aspect	Nature of Questions
Character preference	Character liked most and least, most like to be friends with; Prime character - who child would choose to be
Character Attributes	Realism and smoothness of movement Clothes appreciation and similarity to own age
Character conversations	Content believability, interest and similarity to own conversations
Impact	Victims acceptance of advice and how much child had helped
Bullying Storyline	Storyline believability and length
Similarity	Character that looks and behaves most and least like you
Empathy towards characters	Feeling sorry for characters; if yes which character Feeling angry towards the characters; if yes which character Ideomotoric empathy based on expected behaviour

Table 1: Content of the Agent Evaluation Questionnaire

Results of the evaluation of the scripted version of FearNot! indicated that levels of affect and empathic engagement were higher if children felt that their interactions had an impact on the characters' behaviour [16]. It was notable from responses on the AEQ and from the focus group discussions that a number of children were annoyed that the victim character did not follow their advice, in particular if they tried the same coping response twice, and the victim *still* did not heed the advice, "I didn't think that the characters listened to my advice, as it didn't work". The same pattern of findings emerged for the level of empathy expressed by children towards the characters, with more empathy being expressed towards the characters if they felt that it had followed their coping strategy advice, and if they felt like they had actually helped the victim characters. These results indicate that empathic engagement and the affect of that engagement are increased if the child believes that their presence and interventions has had an impact on the characters.

The emergent narrative approach of this paper is one solution to this problem. A branching structure both lacks flexibility, and in general is subject to unfavourable combinatorics: the 7 coping responses over an eventual 5 or 6 episodes would produce a very large tree, especially given that the outcome in a scenario should depend on the affective states of the characters and the uncertain outcome of physical actions. For instance, if a child advises a victim to hit the bully back, many different narrative outcomes are possible. Depending on its emotional state and its level of confidence, which reflect what has already happened, the victim could decide to deny or follow the advice. Its reaction to the advice would in turn affect the emotional state of the bully, increasing or diminishing its own level of confidence and potentially, altering its action decision too.

A small-scale evaluation was carried out with eleven children randomly chosen from the third and fourth grade in a Portuguese school, making sure that they had no previous contact with the system or the project, and set up as for the scripted evaluation. 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 respond to the victim's request for help by typing free-text. Though the victim had asked for help, the children did not always realize that they could really write something and in this case were prompted by a researcher. 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, which was completed by all subjects, the child completed the same AEQ that had been used in the evaluation of the scripted version. One question was added relating to the dialogue between child and victim: this could not have been used with the scripted version given dialogue had been handled by menu selection. This asks the child if the victim understood the conversation (by giving appropriate responses to the child's inputs) in order to evaluate the competence of this system.

While the scripted version included some recorded speech, 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 dialogues 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 4.

The first two questions refer to the conversation and dialogue between the characters.

Since the episodes displayed are physical bullying episodes, which contain few dialogue lines, and the dialogues in the emergent version are similar in nature 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 in the next episode than to behave as in the first episode.

	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

Table 4. Responses to questions about character responsiveness

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 architecture 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.

Conclusions

A project covering as much ground as this one draws on a very large body of other work in a number of different disciplines, of which we have space here only to reference the most significant. We drew on earlier work using an OCC approach for synthetic agents [13], while the emotion-driven planner was based on [10]. Like most other researchers in this field we must also acknowledge the seminal work of the OZ project [17] and in particular its emphasis on character believability, which was seen as vital to the development of empathy between child and victim.

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 [20] 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 [21] is a beautifully designed story environment, but its conception of beats is close to that of universal plans and produces a very large authoring task that is unlikely to be sustainable for an educational environment. FearNot! and Façade differ in their narrative approaches, the stories in the former being created from, rather than articulated around, the user actions, as it is the case in Façade. The closest work in pedagogical intent is Carmen's Bright IDEAS [8] aimed at teaching cognitive behavioural therapy to mothers of young cancer patients. However this is a dialogue-based application in 2D with a branching structure driven by user selection of one of three 'thought bubbles' at intervals, so that its dramatic form and interactional structure are very different.

As argued above, 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.

Despite the fact that its potential is theoretically demonstrable and that the research community as made a lot of progress with it in the last few years, the vision of an emergent narrative system is one that will only be recognized once computationally implemented. Although it poses many problems, technical and theoretical, it appears from our research that one of the main challenges is the interdisciplinary skills it requires. Interactive dramas based on a bottom up approach cannot be partially implemented. While the graphical design of appropriate characters is a complex task and requires a level of expertise not often available in computer and science laboratories, state of the art agent approaches such as affective architectures, continuous planning and multi-agent interaction models, usually absent from games companies, are also needed. In addition, new testing and evaluation methodologies are required. Many new issues – for example action synchronization and validation, intelligent camera and others – have been raised and more extensive evaluation is required, but the project has met its objective of an emergent and unscripted narrative for anti-bullying education.

References

[1] Aylett, R. 1999. Narrative in virtual environments – towards emergent narrative; AAAI Symposium on Narrative Intelligence pp 83-86

[2] Aristotle, 330 B.C. – 1987. The Poetics of Aristotle, translation and commentary by Stephen Halliwell; Duckworth.

[3] Propp V, 1928 – 1998. Morphology of the Folktale. University of Texas Press, 1998.

[4] Barthes, Roland. 1966. Communication 8, Introduction à l'analyse structurale des récits. Editions du Seuil. 1981

[5] Boal, A. 1999. Legislative theatre: Using performance to make politics. Routledge. 1999

[6] A. Ortony, G. Clore, and A. Collins. The Cognitive Structure of Emotions. Cambridge University Press, New York, reprinted 1994 edition, 1988.

[7] Picard, R. 1997 Affective Computing. MIT Press.

[8] Elliot C.: "The Affective Reasoner: A process model of emotions in a multi-agent system". Northwestern University, PhD Thesis, Illinois, 92.

[9] Fikes, Nilsson. STRIPS: A New Approach to the Application of. Theorem Proving to Problem Solving

[10] Gratch.J (2000) Emile: Marshalling passions in training and education. In Autonomous Agents 2000. ACM Press, 2000.

[11] Louchart,S; Romano,D and Aylett,R.S (2004) Speaking and acting-Interacting language and action for an expressive character. Eds L Sheremetov & M Alvarado, Proc, Workshops on Intelligent Computing, MICAI 2004,. pp308-315

[12] S. Louchart, R. Aylett (2003) Solving the narrative paradox in VEs - lessons from Proceedings, IVA2003 eds T. Rist, R.Aylett, D.Ballin, LNAI 2792 Springer 2003 pp244-248

[13] Martinho, C and Paiva.A (1999) Underwater love: Building tristao and isolda's personalities. In M. Wooldridge and M. Veloso, editors, Artificial Intelligence Today. Springer, 1999.

[14] Russel, S., Norvig, P. (2002)Artificial Intelligence: A Modern Approach. 2ª Edition Englewood-Cliff, N.J.: Prentice-Hall, Ch 12

[15] G. A. Kaminka; M. M. Veloso, S. Schaffer, C. Sollitto, R. Adobbati, A. N. Marshal, A. Scholer, and S. Tejada.(2002) GameBots: The ever-challenging multi-agent research test-bed. Communications of the ACM, Jan 2002.

[16] Hall, L., Woods, S., Aylett, R., Newall, L., & Paiva, A. (2005). Achieving empathic engagement through affective interaction with synthetic characters. Proc ACII 2005, Bejing, China, Springer LNCS 3784, pp. 731-731

[17] Bates, J. (1994). The Role of Emotion in Believable Agents. Communications of the ACM, 37(7) &P 122-125.

[18] Marsella,S and Gratch.J. (2003) Modeling coping behavior in virtual humans: don't worry, be happy. In AAMAS 2003, pages 313–320. ACM Press, 2003.

[19] Ross, D.H; McBriar, D.I; Duke, P. (1998) Rapid Knowledge Elicitation and Expression Graphically in G2. GUS'98—Gensym User Society meeting, May 1998.

[20] Cavazza, M; Charles, F; Mead, S.J (2002): Character-Based Interactive Storytelling. IEEE Intelligent Systems 17(4): 17-24

[21] Mateas,M; Stern, A (2002): A Behavior Language for Story-Based Believable Agents. IEEE Intelligent Systems 17(4): 39-47 (2002)

[22] Aylett, R.S; Dias, J. and Paiva, A (2005) An affectively-driven planner for synthetic characters. Proceedings, UK PlanSIG, City University, to appear Dec 2005

VICTEC: Work partially supported by European Community (EC) under the IST RTD programme, contract IST-2001-33310-victec. The authors are solely responsible for the content of this publication. It does not represent the opinion of the EC, and the EC is not responsible for any use that might be made of data appearing therein.

Dr Lynne Hall is a Principal Lecturer in Interactive Digital Media at the University of Sunderland. She gained her doctoral degree in 1993 at UMIST, UK and then worked in banking before becoming a lecturer. Her main research interests are affective interaction and synthetic characters and empathic engagement. She has over 30 publications in peer-reviewed journals and conferences.
Lynne.hall@sunderland.ac.uk

Sarah Woods (BSc, MSc) is a Lecturer in the School of Psychology at the University of Hertfordshire. She is involved in the E-Circus project, which involves educational virtual dramas about bullying behaviour and the social challenges faced by refugees, for primary school children. Previously, she worked at the Institute of Psychiatry, London (2000-2002) and the University of Hertfordshire (1998-2000).
s.n.woods@herts.ac.uk

Sandy Louchart (BSc, MSc) is a research associate in the School of Maths and Computer Science at Heriot-Watt University. He is involved on the European funded Project, E-Circus. Prior to this, he has worked on the VICTEC (Virtual ICT with Empathic Characters) project (2002-2005). He researches actively the domains of Intelligent Agents, Interactive Narratives, Serious Games and emotion-based agent modelling.
sandy@macs.hw.ac.uk

Ruth Aylett is Professor of Computer Sciences in the School of Maths and Computer Science at Heriot-Watt University. She researches the fields of Intelligent Virtual Environments, Synthetic characters and narrative. She managed the VICTEC project between 2002 and 2005 and is now leading the Ecircus project. Prior to this position, she worked at Sheffield Hallam University, Edinburgh University and Salford University.
ruth@macs.hw.ac.uk

Ana Paiva is a research group leader at INESC-ID and a Professor at Instituto Superior Técnico, Technical University of Lisbon. Well known in the area of Intelligent Agents, User Modelling and AI Applied to Education, her research focuses on the affective elements in the interactions between users and computers (synthetic characters). She has (co)authored over 80 publications in refereed journals, conferences and books.
ana.paiva@inesc-id.pt

Marco Vala (BSc, MSc) is a researcher at the Intelligent Agents and Synthetic Characters Group (GAIPS) of INESC-ID, and a teaching assistant of information systems and computer engineering at IST, Technical University of Lisbon. He participated in the EU-funded projects SAFIRA, ELVIS and VICTEC. Currently undertaking a PhD, his research focuses on embodied characters, human-character interaction, and intelligent bodily behaviour.
marco.vala@tagus.ist.utl.pt.

Joao Dias (BSc, MSc) is a researcher at the Intelligent Agents and Synthetic Characters Group (GAIPS) of INESC-ID, and a teaching assistant of information systems and computer engineering at IST, Technical University of Lisbon.
Joao.dias@tagus.ist.utl.pt.