Unscripted Narrative for Affectively Driven Characters

Ruth Aylett and Sandy Louchart
Heriot-Watt University

Joao Dias, Ana Paiva, and Marco Vala
INESC-ID

Sarah Woods
University Hertfordshire

Lynne Hall
University Sunderland

The authors discuss designing unscripted (emergent) dramas with affectively driven intelligent autonomous characters based on the development of the FearNot! system for education against bullying.

The video game industry has successfully demonstrated over the last two decades that virtual characters, virtual worlds/environments, and even virtual societies can reach and entertain large numbers of users. However, with some exceptions (for example, *The Sims* and *Black and White*) much of the development effort over this period has gone into improving graphics quality 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.

But the emergence of educational applications and the developing interest of the education community in the use of VR technologies have raised important issues of narrative articulation and story representation in virtual environments. Because 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 (such as role-playing, storytelling, and participative activities).

In this article, we report on the emergent narrative concept[1] aiming at the definition of a narrative theory adapted to the VR medium (whether a 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 doesn't seem possible to resolve it through the use of existing narrative theories. Interactivity is the novel element that storytellers must address: from Plato's story definition[3] and Aristotle's plot consideration all the way to Propp's metastructural narrative articulation;[4] Campbell's cyclical diagrams;[5] and Barthes poststructuralist, top-down analytical masterpiece.[6]

A form in which the audience is not static and can interact with the story's characters or environments brings another dimension to storytelling and extends the boundaries of both narrative creation and articulation. A story can from now on not only be told by the author/poet directly (diegesis, in Plato's terminology[3]) or shown to the audience through the use of characters (mimesis), but also be experienced and lived through.

Keeping this in mind, we decided to try an emergent-narrative approach in developing our Fun with Empathic Agents Reaching Novel Outcomes in Teaching (FearNot!) application. The primary goal of our system was the creation of empathetic characters that would help educate children against bullying. (To learn more on why we chose this approach, see the sidebar, "Choosing an Emergent-Narrative Solution.")

Choosing an 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 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

Published by the IEEE Computer Society

Table A. Examples of character and drama triggers and events.

Context	Trigger Type	Event Type
Character Management		
The player isn't interacting/attentive	No interaction when there's an opportunity	Send a nonplayer character (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 player of roles and rules, bring in next encounter
Player incorrectly determines what to do next	Player pursues wrong goal, goes in the wrong direction	Give hints that the player is going the wrong way, or have empty screen

articulation, we've 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 theater such as Boal's Forum Theatre and street theater,[1] reality TV, the role of nonplayer characters in video games, and in particular role-playing games (RPGs) in their many forms and aspects (such as cooperating and conflicting board-based or live games).

Research into RPGs is mainly empirical since some practitioner literature exists, but few scholarly resources are available for detailed study. Using the approach of knowledge engineering, we interviewed 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 parties have conflicting goals), and one in live RPGs—using a professional knowledge acquisition tool (KAT) system: KAT Builder.[2] We used the KA tool to extract story management rules, discussed in more detail elsewhere[3]—we show a subset in Table A.

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. This replaces the artifact-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 center of the story. Thus a character-based approach is not composed of a single storyline to which the different characters must conform, but of as many storylines as there are characters.

It's 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, because each character is in the center of his or her narrative, the focus of development in an actual implementation must be oriented toward the completeness of the character.

Emergent narrative does not abolish authoring: instead, 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 among themselves, their different potential relationships with each other must be thought through and their place in the world must be clearly established.

In addition to the creation of both worlds and characters that would likely interact in interesting and potentially dramatic ways, the author must formulate the setup 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's 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 modeled.

Although a rich definition of character and modeling 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. Because the character is at the center of narrative development from its own and the system's perspective, the development of intelligent agents that can react and therefore act autonomously under certain stimuli (narrative, emotional, and personality related) are a requirement for translating theory into implementation.

References

1. A. Ortony, G. Clore, and A. Collins, *The Cognitive Structure of Emotions*, Cambridge Univ. Press, 1994.
2. D.H. Ross, D.I. McBriar, and P. Duke, "Rapid Knowledge Elicitation and Expression Graphically in G2," *Proc. Gensym User Soc.* (GUS), GUS, 1998.
3. S. Louchart and R. Aylett, "Solving the Narrative Paradox in VEs—Lessons from Proceedings," *Proc. Intelligent Virtual Agents* (IVA), T. Rist, R. Aylett, and D. Ballin, eds., Springer, 2003, pp. 244-248.

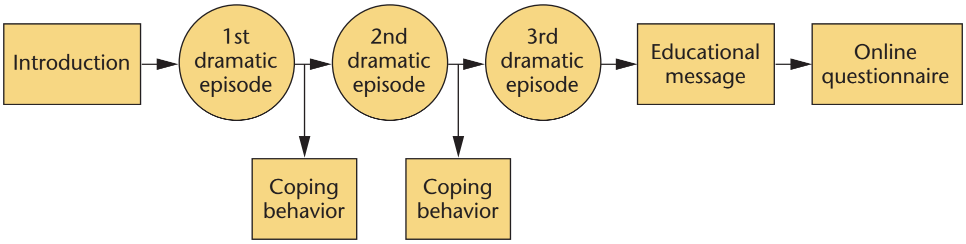

1 Interactive structure of FearNot!

2 A bullying incident.

Application

Affect is seen as central to the creation of unscripted narrative, since it both produces dramatically interesting action selection and the accompanying expressive behavior required to establish an action's context in a character's motivations. In the rest of this article, we discuss an initial experiment in implementing emergent narrative, carried out in the European Union Framework V project, Virtual Information and Communication Technology (ICT) with Empathic Characters (VICTEC).

The agent framework we developed for the project allows the construction of virtual intelligent agents that express and react to emotions in a natural and meaningful way. We designed it so that it doesn't only apply to the specific context of school bullying—we can also use it in the more general realization of emergent dramas.

VICTEC, involving five partners in the UK, Germany, and Portugal, sought to apply virtual dramas acted out by 3D graphically embodied characters to what's known generically in the UK as Personal and Social Education (PSE). This covers topics such as education against bullying and racism, about drugs (including smoking and alcohol), and related to sex.

A common thread in these topics is that knowledge in and of itself isn't sufficient to meet the pedagogical objectives, since attitudes and emotions are at least as important in producing a desired behavior. For this reason, techniques such as small-group discussion, role-play, and dramatic performance by theater-in-education groups can be used.

The project's aim was to create some of the impact of dramatic performance through virtual dramas. The group selected antibullying education as its focus. Effective though theater-in-education is in this domain, it's necessarily collective, and in any group it's likely that some individuals will be victims of bullying in real life and therefore would feel inhibited about participating. This led us to suggest a virtual drama application that individuals could use.

The aim of the FearNot! 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. We felt that the creation of an empathic relationship between the child and autonomous character could act as the mechanism through which this sense of responsibility would be achieved, so that the child user would really care about what happened to the victimized character. The child was asked to act as an invisible friend, and to give advice that would influence the behavior 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[7] to incorporate theater 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 with the actor, who stays in role and negotiates with the group about what he or she should do next in the drama, respecting the constraints of the role and character. We incorporated this structure of dramatic episodes divided by periods in which advice can be given to a character for FearNot! as Figure 1 shows schematically.

The session starts with an introduction to the school and the characters and then the first dramatic episode follows, 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 behavior. This structure is repeated—currently twice, but with a target of five or six episodes—and a simple educational message is displayed, followed by an online questionnaire assessing how far the child can relate to the experiences of the characters just seen.

Agent framework

Figure 3 shows the agent architecture used in the FearNot! demonstrator. Because character interaction is what drives the events during an episode, the appraisal-driven agent architecture forms a central part of the system.

With an emergent narrative mechanism, it's the character's ability to autonomously decide upon its own actions—its 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, and so on) and acts on the environment through its effectors. This allows the character to perform different actions (for example, a bully may hit the 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' behaviors, 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 for a more complex and rich behavior.

The appraisal process feeds the resulting emotional state into action selection at two levels: action tendencies and coping behavior. For example, if the victim character starts to cry when bullied, it's not because she or he has a goal that involves crying—this is an innate reaction to a particular distressed emotional state and the inability to fight back.

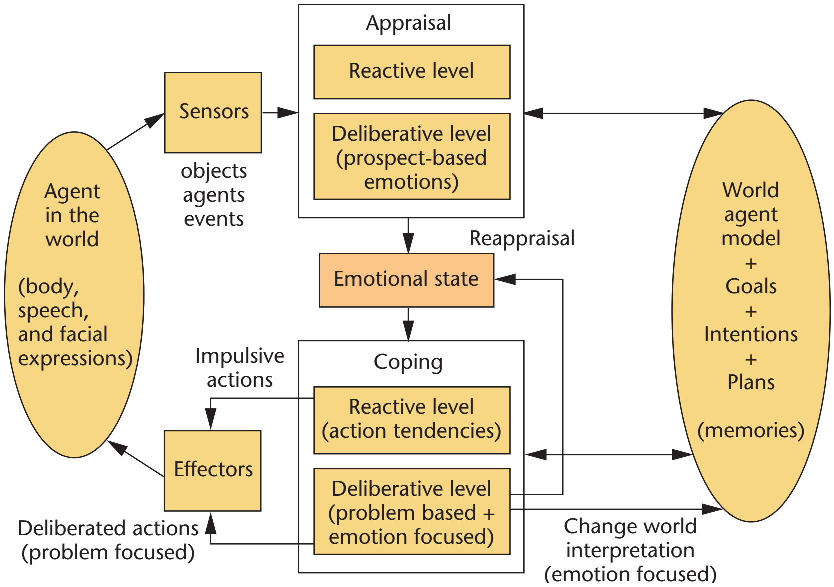

3 **Affective agent architecture.**

Table 1. FearNot! character emotion attributes.

Attribute	Description
Type	Type of emotion being experienced
Valence	The basic types of emotional responses (positive or negative)
Target	Name of the agent/object toward which the emotion is directed
Cause	Event/action that caused the emotion
Intensity	Intensity of emotion; a logarithmic scale between 0 and 10
Time stamp	Moment in time when the emotion was created

Emotion model

The emotion definition we adopted for FearNot! is that of Ortony, Clore, and Collins (OCC).[8] The OCC model is an approach based on appraisal of the an event's affective valence (good or bad) and intensity of impact; the classification of emotions the event defines can be seen as a hierarchical taxonomy of 22 emotion types. Using this model, we represented emotions in our architecture by attributes (see Table 1).

Each emotion type can be realized in a variety of related forms with varying degrees of intensity (for example, the emotion type "fear" can generate an emotion range from concerned 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 the emotional impact and potential answer to the stimulus.

As a dynamic process, an emotion's intensity must be attenuated through time from the moment it's generated onward. We modeled this using Picard's[9] decay function for emotions, which characterizes intensity as a function of time. At any time (t), the value for the intensity of an emotion (em) is given by the formula:

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

The value d (decay) determines how fast the this particular emotion's intensity 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 for the following:

- 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, and
- the character's action tendencies.

Ortony et al.[8] 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 aren't 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:

$$\text{Intensity}(em) = \text{Potential}(em) - \text{EmotionThreshold}(em)$$

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

4 Operator example.

Thus a calm character will have a high threshold and a strong decay for the "anger" emotion type, and will rarely experience anger. When the calm character does experience anger, it's with lower intensities and for a shorter period of time. We should also note that it's possible to have two characters with the same goals, standards, and behaviors that react differently to the same event (by having different thresholds).

Appraisal mechanism

As Figure 1 shows, we made our appraisal mechanism with two distinct layers. The reactive layer appraisal is handled by a set of emotional reaction rules, based on Elliot's Construal Theory.[10] 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, and so on).

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 those that the character actively tries to achieve (such as going to a dental appointment) while interest goals represent goals that the character has but doesn't pursue (such as avoiding getting hurt). We did not use the OCC replenishment goals since they could be considered as active-pursuit goals with cyclic activation and deactivation.

When the character appraises an event, the deliberative level checks if any goal has become active, and if so, creates an intention to achieve the goal's success conditions, 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, changing plan probabilities, and generating new hope and fear emotions.

Action selection and coping mechanisms

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 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 (for instance, triggered by the same emotion), the most specific one is chosen.

To build the core of the deliberative or coping layer we implemented a standard partial-order continuous planner[11] and extended it to include emotion-focused coping strategies in addition to the common planning operations.[12] 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 might be generated for this target goal and the planner selects one to continue planning or execution. Once the selected plan is brought into focus it generates hope and 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 modeled as a set of operators along with ordering constraints, causal links, binding constraints, and open preconditions. Operators are a slight modification of STRIPS[13] operators, associating probability values with the effects. Figure 4 shows the operator for the GetUp action. To get up, the agent can't 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 can get up if it's seated or lying down.

Emotion-focused coping

Marsella and Gratch[14] 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 can't be achieved, so that failure is accepted. When a plan has a low probability of success, the planner will accept the plan's failure and won't 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—that is, 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 doesn't have a high probability of happening, the agent can ignore the threat assuming that the effect will never happen by lowering the effect probability.

This agent architecture, 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 alter the agent's emotional state.

Since the agent makes decisions based on that emotional state, this potentially affects its perception of actions, alters the probability of the plan's success, and affects 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 (according to their context) without the need for scripting them.

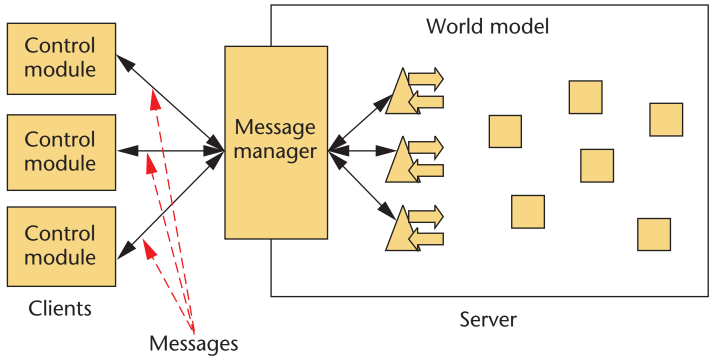

5 Client–server architecture.

Visualization system

Figure 1 shows only a part of the FearNot! architecture: the agent mind that controls character behavior and the virtual world model to which it's linked by sensors and effectors. Underlying this is a client–server architecture (see 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 initializing episodes), and a view manager all operate around the symbolically represented central world model (see Figure 6).

A visualization system presents this virtual world that the characters inhabit to the user. The view manager—a special agent (which doesn't represent a character) with the power to listen to all the events that occur in the virtual world—translates those events to a specific visualization system, which in the case of FearNot! is 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 specify the way the character will be visualized—for example, the character's facial expression—as well as map planner actions onto animations defined in the visualization system.

This modular separation of artificial intelligence and graphics components makes each editable or replaceable: WildTangent could be replaced by, for example, the *Unreal Tournament* engine, 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 different from that commonly employed by graphics researchers, in which the whole architecture is embedded in the visualization system. For example, with *Unreal Tournament*, scripting represents an extension and generalization of work such as GameBots,[15] in which agent minds are directly attached to a visualization engine via sockets. In our work an added virtual world model is the central and key component of the architecture and message; passing is generalized.

Of course, the character bodies and animations, as well as the 3D graphic visualization of the rest of the virtual world, have to be specified in a format that the visualization engine will accept and load at runtime. The visualization engine also has to deal with specific issues relating to 3D space, including path planning of the movements specified by the character minds. Additionally, in FearNot!, the outcome of actions that aren't certain to succeed are also decided in the visualization engine—for example, whether a character who's 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

In this section, we examine an example of an emergent narrative—the scenario used in the evaluation discussed next—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 dislikes John and when he sees John he starts insulting him (reactive action tendency). As a result, John has an active pursuit goal of fighting back that's triggered when he's insulted by other characters. He tries to build a plan 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

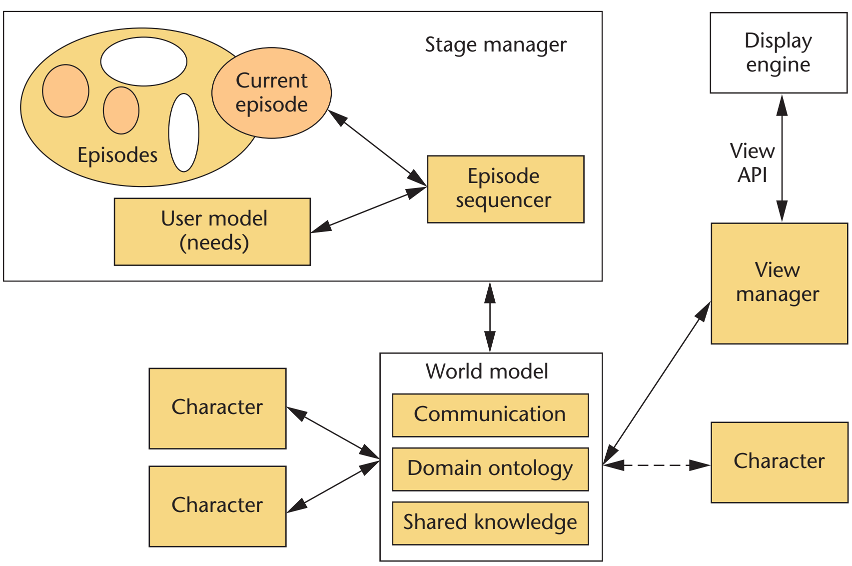

6 FearNot! overall architecture.

7 **John (the victim) and Luke (the bully) in the classroom.**

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 doesn't 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 situation. Luke feels confident of his plan, so he starts walking toward the book with a happy face (the hope emotion is mapped to a happy facial expression). On the other hand, John feels distressed at being insulted and disappointed by not being able to fight back. Luke moves toward 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 toward 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's decided by physics, and in the virtual world by a probability set in the 3D visualization—a character might or might not fall. If John falls, he appraises this event as undesirable and activates an action tendency to start crying. At the same time, Luke appraises the same event as desirable and starts gloating by saying something like "What a wimp, I hardly touched you." When John cries, Luke finds it blameworthy and thus threatens him to stop crying and to not tell anyone. If John doesn't fall, Luke won't mock him. Instead, John might feel angry and ask Luke why he's always picking on him. Luke responds negatively to the question by insulting John even more.

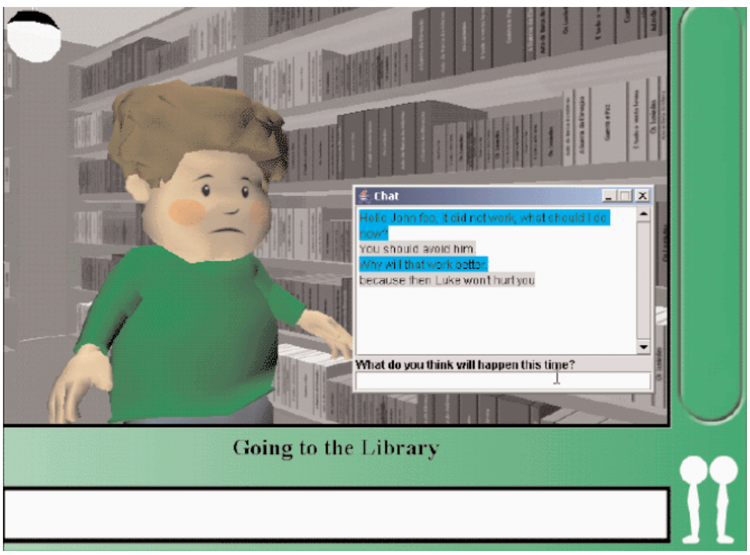

8 **Child interacting with the victim.**

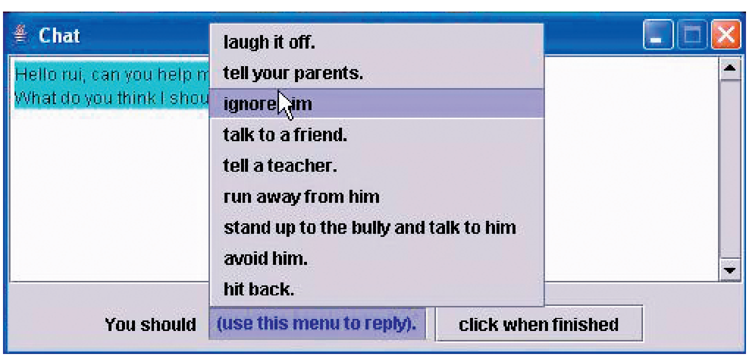

9 **Menu of coping choices in the interface.**

Figure 8 shows a snapshot of the interaction mode in which the child user talks with the character victim and advises him or her on what to do next. The user types whatever she or he wants in the lower text box on the right and presses the OK button, sending the written utterance 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.[16]

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 doesn't distinguish 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 isn't strong enough to make him give up the goal. Thus, user interaction changes the victim's behavior by indirect influence rather than because the victim does exactly what he's told. However, if John tries pushing Luke and it doesn't succeed, then he won't accept a further suggestion to hit back, since the experience of being hurt again alters his emotional state, this time in the direction of greater fearfulness.

Evaluating FearNot!

We carried out a large-scale evaluation (using 401 children) of a scripted version of FearNot! at the University of Hertfordshire in 2004. This took place in a large information technology suite for a period of two weeks, with up to 65 third- and fourth-grade children participating each day.[17] Our aim was to evaluate the overall pedagogical concept and the scenario material used before adding the autonomous agent architecture. We seated children individually at workstations and the project coordinator provided them with 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 similar to the one shown in Figure 1. In this version, the child selected a piece of advice from a randomly ordered list of coping strategies (shown as a drop down menu—see Figure 9), rather than entering free text as in the later unscripted version. This dialogue 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, the application displayed a universal educational message noting 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 wasn't influenced by or dependent on the child's choice of coping strategies.

Table 2. Content of the agent evaluation questionnaire.

Aspect	Nature of Questions
Character preference	Which character they liked most and least, which they would like most to be friends with, and who was the prime character (who would the child 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	Victim's 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 toward characters	Either feeling sorry for a character (if yes, which character), feeling angry toward a character (if yes, which character), and ideomotoric empathy based on expected behavior

After children interacted with FearNot!, they completed an agent evaluation questionnaire (AEQ), and participated in a focus group. Table 2 shows the questionnaire's structure. Children's views were predominantly measured according to a five-point Likert scale. We systematically developed the AEQ throughout the course of the project and took extreme care to ensure that children understood the terminology and language used as well the rating scale. We carried out several pilot studies using the AEQ prior to the large-scale evaluation event.

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' behavior.[17] 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 didn't follow their advice, in particular if they tried the same coping response twice, and the victim still didn't heed their advice. We commonly heard feedback such as, "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 toward the characters, with more empathy being expressed toward a character if they felt that the character followed their coping strategy advice, and if they felt that they had actually helped the victim character. These results indicate that empathic engagement and the effect of that engagement increase if the children believe that their presence and interventions have an impact on the characters.

The emergent-narrative approach presented in this article is one solution to this problem. A branching structure lacks flexibility, and in general is subject to combinatorial explosion: the seven coping responses over an eventual five or six episodes would produce a 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 the character's emotional state and confidence level, 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.

We carried out a small-scale evaluation of the emergent-narrative approach with 11 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. We used the physical bullying story just described and each child participated individually. After the initial introduction and the first episode, we asked each child to respond to the victim's request for help by typing in their advice. Although the victim had asked for help, the children didn't always realize that they could really write something and in this case a researcher prompted them. 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 we used in the evaluation of the scripted version. We added one question relating to the dialogue between the child and victim (this couldn't have been used with the scripted version, since dialogue was handled through a menu selection). This additional question asked the child if the victim understood the conversation (by giving appropriate responses to the child's inputs) to evaluate the competence of the system.

While the scripted version included some recorded speech, the emergent version has no sound at all. This is

Table 3. Responses to questions about character responsiveness.

Questions	Scripted Version	Emergent Version
Did the conversations seem real? (yes = 1; no = 5)	2.4	1.9
How 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

a disadvantage, as the episodes might not seem so engaging, making story comprehension 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, 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 (see Table 3). 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 strategy to fight back, it seems more real to see him threatening the bully in the next episode than to behave as in the first episode.

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 narrative development, the whole agent architecture and the surrounding framework allowing agents to interact with each other have to be completely in place before we can carry out any real testing of the narrative produced. This is different from a top-down approach, in which a subset of facilities can typically be made available early and then elaborated. In particular, if we are to present emergent narrative graphically, the graphic visualization must support full-agent autonomy, including movement in the environment and the execution of animations. Because of the way we designed the graphical world 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 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. For the system to 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 characters' actions and movement 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 must also be increased.

Speech output raises particular problems, too, in an unscripted environment. The template-based language system developed for FearNot! seems perfectly capable of generating the range of utterances needed for intercharacter dialogue. The system also coped rather better than we hoped in both English and Portuguese, with character–child dialogue. However, given the robotic nature of text-to-speech synthesis systems, we decided at an early stage to stick to text output on the screen rather than destroy the characters' believability. 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 large database, which was incompatible with the resources available when running interactive graphics. Moreover, those types of systems have been designed for adult voices only with the equable tones of a telephone help system, not the angry or miserable child characters of FearNot!

The team considered a methodological point about the use of this approach in an educational application. To what extent is the somewhat unpredictable outcome of episodes in conflict with our pedagogical objectives? It's possible, for example, for the stage manager to bring characters together so that bullying will take place and then for nothing 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 might be that the amount of narrative variability that's acceptable will depend on the exact application chosen.

Final remarks

As we previously argued, many issues have arisen from the emergent-narrative work carried out in FearNot! and further research is required to deal with these. (For a brief overview of other narrative work being developed, see the "Related Work" sidebar.) How-

Related Work

A project covering as much ground as this one draws on a 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[1] using the Ortony, Clore, and Collins approach for synthetic agents, while we based our emotion-driven planner on Gratch's work.[2] Like most other researchers in this field, we must also acknowledge the seminal work of the OZ project,[3] in particular its emphasis on character believability, which was seen as vital to the development of empathy between the child and victim.

However, apart from the novelty of the application domain—no previous autonomous agent application has targeted antibullying education—the emergent-narrative experiment was also truly novel in our view. Much other interesting storytelling work exists, but no other group seems as yet to have attempted an unscripted approach in this way. Variation in the story outcome has been generated for example by Cavazza et al.,[4] but this is derived from prebuilt 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[5] is a beautifully designed story environment, but its conception of beats is close to that of universal plans and produces a large authoring task that isn't likely 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 is the case in *Façade*. The closest work in pedagogical intent is Carmen's Bright IDEAS[6], aimed at teaching cognitive behavioral 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 interactive structure are quite different.

References

1. C. Martinho and A. Paiva, "Underwater Love: Building Tristao and Isolda's Personalities," *Artificial Intelligence Today*, M. Wooldridge and M. Veloso, eds., Springer, 1999.
2. J. Gratch, "Emile: Marshalling Passions in Training and Education," *Proc. Autonomous Agents*, ACM Press, 2000.
3. J. Bates, "The Role of Emotion in Believable Agents," *Comm. ACM*, vol. 37, no. 7, 1994, pp. 122-125.
4. M. Cavazza, F. Charles, and S.J. Mead, "Character-Based Interactive Storytelling," *IEEE Intelligent Systems,* vol. 17, no. 4, 2002, pp. 17-24.
5. M. Mateas and A. Stern, "A Behavior Language for Story-Based Believable Agents," *IEEE Intelligent Systems*, vol. 17, no. 4, 2002, pp. 39-47.
6. S. Marsella, "Pedagogical Soap," *Socially Intelligent Agents: The Human in the Loop,* AAAI Press, 2000.

ever, we believe we've shown that there's an interesting role for this approach to unscripted narrative, and that applications with an open-ended and somewhat unpredictable narrative could have much to offer.

Despite the fact that its potential is theoretically demonstrable and that the research community has 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 technical and theoretical problems, 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 can't 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 multiagent interaction models—usually absent from games companies—are also needed. In addition, new testing and evaluation methodologies are required. Issues such as action synchronization and validation, intelligent cameras, and others need to be addressed. More extensive evaluation is required, but the project has met its objective of an emergent and unscripted narrative for antibullying education. ■

Acknowledgments

Our VICTEC work was partially supported by the European Community (EC) under the IST RTD program, 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 the data appearing therein.

References

1. R. Aylett, "Narrative in Virtual Environments—Towards Emergent Narrative," *Proc. Am. Assoc. for Artificial Intelligence* (AAAI) *Symp. Narrative Intelligence*, 1999, AAAI Press, pp. 83-86.
2. S. Halliwell, *The Poetics of Aristotle: Translation and Commentary*, Univ. of North Carolina Press, 1987.
3. Plato, *La Republique*, GF-Flammarion, 1966.
4. V. Propp, *Morphology of the Folktale*, Univ. of Texas Press, 1998.
5. J. Campbell, *The Hero with a Thousand Faces*, Princeton Univ. Press, 1949.
6. R. Barthes, "Introduction à l'Analyse Structurale des Récits" [Introduction to the Structural Analysis of the Accounts], *Communication 8,* 1966 (in French).
7. A. Boal, *Legislative Theatre: Using Performance to Make Politics*, Routledge, 1999.
8. A. Ortony, G. Clore, and A. Collins, *The Cognitive Structure of Emotions*, Cambridge Univ. Press, 1994.
9. R. Picard, *Affective Computing*, MIT Press, 1997.
10. C. Elliot, "The Affective Reasoner: A Process Model of Emotions in a Multi-Agent System," doctoral dissertation, Northwestern Univ., 1992.
11. S. Russel and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed., Prentice-Hall, 2002, ch. 12.

12. R.S. Aylett, J. Dias, and A. Paiva, "An Affectively Driven Planner for Synthetic Characters," *Proc. Int'l Conf. Automated Planning and Scheduling,* 2006; http://www.plg.inf.uc3m.es/icaps06/acceptedpapers.htm.
13. R. Finkes and N. Nelson, *STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving,* vols. 3-4, no. 2, 1971, pp. 189-208.
14. S. Marsella and J. Gratch, "Modeling Coping Behavior in Virtual Humans: Don't Worry, Be Happy," *Proc. Adaptive Agents and Multi-Adaptive Systems* (AAMAS) *2003,* ACM Press, 2003, pp. 313-320.
15. G.A. Kaminka et al., "GameBots: The Ever-Challenging Multi-Agent Research Test-Bed," *Comm. ACM*, Jan. 2002, pp. 43-45.
16. S. Louchart, D. Romano, and R.S. Aylett, "Speaking and Acting—Interacting Language and Action for an Expressive Character," *Proc. Workshops Intelligent Computing*, L. Sheremetov and M. Alvarado, eds., Sociedad Mexicana de Inteligencia Artificial, pp. 308-315.
17. L. Hall et al., "Achieving Empathic Engagement through Affective Interaction with Synthetic Characters," *Proc. Affective Computing and Intelligent Interfaces* (ACII) *2005*, Springer, 2005, pp. 731-731.

Ruth Aylett *is a computer science professor in the School of Maths and Computer Science at Heriot-Watt University. Her research interests include intelligent virtual environments, synthetic characters, and narrative. Aylett has a BSc from the University of London and is a chartered engineer. Contact her at ruth@macs.hw.ac.uk.*

Sandy Louchart *is a research associate in the School of Maths and Computer Science at Heriot-Watt University. His research interests include intelligent agents, interactive narratives, serious games, and emotion-based agent modeling. Louchart is pursuing a PhD in interactive narrative from the University of Salford. Contact him at sandy@macs.hw.ac.uk.*

Joao Dias *is a researcher at the Intelligent Agents and Synthetic Characters Group of INESC-ID, and a teaching assistant of information systems and computer engineering at Instituto Superior Técnico (IST), the Technical University of Lisbon. His research interests include affective agent architecture and emotion-based planning. Dias has an MSc in computer science from IST. Contact him at Joao.dias@tagus.ist.utl.pt.*

Ana Paiva *is a research group leader at INESC-ID and a professor at the IST, Technical University of Lisbon. Her research interests include intelligent agents, user modeling, artificial intelligence applied to education, and affective elements in user and computer interactions. Paiva has a PhD from the University of Lancaster in artificial intelligence in education. Contact her at ana.paiva@inesc-id.pt.*

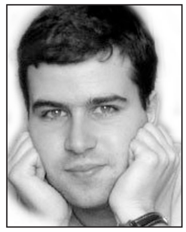

Marco Vala *is a researcher at the Intelligent Agents and Synthetic Characters Group of INESC-ID, and a teaching assistant of information systems and computer engineering at IST, Technical University of Lisbon. His research interests include embodied characters, human– character interaction, and intelligent bodily behavior. Vala is pursuing a PhD in synthetic characters from the IST. Contact him at marco.vala@tagus.ist.utl.pt.*

Sarah Woods *is a lecturer in the School of Psychology at the University of Hertfordshire. Her research interests include educational virtual dramas about bullying behavior for primary school children and the social challenges faced by refugees. She is a member of the British Psychological Society. Contact her at s.n.woods@herts.ac.uk.*

Lynne Hall *is a principal lecturer in interactive digital media at the University of Sunderland. Her research interests include affective interaction and synthetic characters and empathic engagement. Hall has a PhD in user interface design for multi-agent systems from the University of Manchester, UK. Contact her at Lynne.hall@sunderland.ac.uk.*