Feeling and Reasoning: a Computational Model for Emotional Characters

João Dias and Ana Paiva

GAIPS Grupo de Agentes Inteligentes e Personagens Sintéticas,
INESC-ID and IST, Tagus Park
Av. Prof. Cavaco Silva, 2780-990 Porto Salvo, Portugal
Email: joao.assis@tagus.ist.utl.pt, ana.paiva@inesc-id.pt

Abstract. Interactive virtual environments (IVEs) are now seen as an engaging new way by which children learn experimental sciences and other disciplines. These environments are populated by synthetic characters that guide and stimulate the children activities. In order to build such environments, one needs to address the problem of how achieve believable and empathic characters that act autonomously. Inspired by the work of traditional character animators, this paper proposes an architectural model to build autonomous characters where the agent's reasoning and behaviour is influenced by its emotional state and personality. We performed a small case evaluation in order to determine if the characters evoked empathic reactions in the users with positive results.

1 Introduction

The art of creating engaging and believable characters is well studied among traditional animators [23]. Traditional characters like *Mickey Mouse*, or more recent 3D characters like *Shrek*, are able to create the illusion of life and allow for the establishement of emotional relations by the viewers. The viewer feels sad when they are sad, angry when something unfair is done to them and so on. These emotional relations are named empathic relations. Empathy can be defined in broad terms as "an observer reacting emotionally because he perceives that another is experiencing or about to experience an emotion" [7].

The use of such empathic characters in virtual learning environments has obvious advantages. Children's didactic software usually uses animated characters (3D or not) to guide the child trough the application and activities. They stimulate the child interaction with the environment, enrich the child experience and captivate their attention. However, such animated characters, like in traditional animation, are scripted for each possible scenario when the application is designed. When the child presses a given button, the character will just play the corresponding scripted behaviour. This does not only forces to create such scripted animations for each possible situation, but also limits the possibilities of the child's interaction.

Ideally, one would like to have Intelligent Virtual Environments (IVEs) inhabited by autonomous agents, which "think" and act on their own. Such au-

tonomous agents make the environment neither predictable nor completely controlled, and thus it is not possible to prescript animations for each situation. The narrative can then emerge from the individual performance of each character. However, making autonomous agents believable and empathic it's a quite difficult problem. This paper presents an agent architecture that aims at achieving such empathic autonomous characters, inspired by some of the elements present in traditional animation.

To illustrate our approach, we will look at one particular example of a pedagogical system. FearNot![5] is a computer application developed to tackle and eventually help to reduce bullying problems in schools. Bullying has associated with it a wide variety of behaviours such as hitting, or kicking, in the case of direct bullying, or, in relational bullying, social exclusion or malicious rumour spreading. Thus, the overall objective of the development of FearNot!, was to build an anti-bullying demonstrator in which children age 8 to 12 experience a virtual scenario where they can witness (from a third-person perspective) bullying situations. The child acts as an invisible friend to a victimized character, discussing the problems that arise and proposing coping strategies.

Note that in bullying situations there are quite clear identifiable roles: the bully, the victim, bully-victim (a child that is sometimes the victim and sometimes the bully) and bystander. Therefore it is necessary to build an agent architecture that not only supports believability, but also offers an easy process of building characters with particular behaviours. In sum, the architecture aims at achieving synthetic characters with the following characteristics:

- **Believability and Empathy:** The characters must be believable and be able to produce empathic reactions with users.
- **Reactive and Cognitive Capabilities:** Given the scope of possible domains, characters should react as quickly as necessary in a rapidly changing environment. However, reactive behaviour is too predictable for a truly autonomous character. Believable characters should display motivations, goals and desires, which is only possible if they have cognitive capabilities.
- **User Interaction:** The characters should be able to interact with an external user and receive suggestions. However, any influence the user may perform cannot be direct, because the character cannot take the user suggestions blindly without taking the risk of not acting in character and thus loosing believability.
- **Generality:** The agent architecture should be domain independent, i.e. it must allow the easy creation of different characters with different personalities for different domains.

To determine if the developed architecture is able to achieve believable characters a small evaluation was performed with eleven-year old children. In order to determine the effects of the user's interaction in the story, we compared our results with the results obtained from a scripted version, where the children could not influence the outcome of the story.

The rest of the paper is organized as follows: in section two we present some related work that led us to the final design of our architecture; next we define some of the most relevant concepts used in our model and depict the architecture. Afterwards an illustrative example is presented to explain how the internal

mechanisms achieve the overall behaviour; finally we analyse the results obtained from the evaluation and draw some final remarks.

2 Related Work

In order to achieve believable characters, our work focuses mainly on two characteristics early pointed out by traditional animators and often explored by researchers working in synthetic characters: emotional expressivity and personality. Characters that are unable to express their feelings and cannot react emotionally to events are lifeless. As Thomas and Johnston put it: "... it has been the portrayal of emotions that has given the Disney characters the illusion of life" [23]. They define three major requisites to successfully express emotions with characters: (1) the emotional state of the character should be undoubtedly perceived by the audience; (2) emotions can be accentuated or exaggerated to be better understood; and (3) emotions affect the reasoning process and consequences should be noticeable in the actions of the characters.

Together with emotional expressivity, personality plays a very important role in believability. Thomas and Johnston state that if the process of thought reveals the feeling, it should also reveal the character's personality. This means that like the emotions, personality also influences the reasoning process. Two different characters may act differently in the same situation because of their personality, and the viewers must perceive such differences. Furthermore, a well defined role and personality is crucial to achieve emergent narrative.

Realizing the importance of emotions in reasoning, several psychologists developed emotion theories that model the generation of emotions in human beings. One of the most important, especially for the computer science community, is the OCC theory of emotions (named after its creators Ortony, Clore and Collins) because it is an appraisal theory that is easily implemented by a computer program [17]. OCC defines emotions as a valenced (good or bad) reaction to an event, which is triggered by a subjective evaluation of the event according to the character's goals, standards and beliefs.

Several researchers used OCC to explore the role of emotions in behaviour. For instance, Joseph Bates realized the importance of believable characters (a term introduced by him [3]) in Virtual Reality applications [2] and together with Reilly [19] used OCC to model emotions in the Oz project [1]. Elliott was also one of the first ones to use OCC in his Affective Reasoner [8]. Another researcher that has further explored the use of OCC theory, Jonathan Gratch, introduced the concept of emotional planners in the Émile system [11]. Gratch argues that planning algorithms have several properties that ease cognitive appraisal. Instead of using domain specific rules to determine the appraisal, Émile takes advantage of explicitly storing the agent plans into memory to reason about future possible outcomes and to automatically generate the character's emotional state. As plans grow and change, so changes the emotional state.

Looking at a different perspective, emotions can also play a significant role in coping strategies. Usually, characters act in the environment to solve their prob-

lems (problem-focused coping). Stacy Marsella introduced a new coping concept in Carmen Bright's Ideas (CBI) [15]. The characters can change their interpretation of the world to discharge negative emotions and to reinforce positive ones, which enables a mental coping mechanism guided by emotions (emotion-focused coping). For example, a problem-focused way to attempt to deal with a loved one's illness, is to take action that gets them medical attention. On the other hand, emotion-focused strategies may include avoiding thinking about it, focusing on the positive or denying the seriousness of an event. Gratch and Marsella further extended and integrated their ideas to create EMA which is used in the Mission Rehearsal Exercise system [12].

Our proposed architecture uses a multi-layered approach similar to TABASCO architecture [22]. The cognitive layer is mainly inspired by Gratch and Marsella's work. A emotional planner builds up the core of the cognitive layer and emotional-focused coping differentiates behaviours according to the characters' personality. The reactive layer is based on the emotional rules used in Martinho's work [16].

3 Emotion and Dynamics of Emotion

Our concept of emotion steams from OCC cognitive theory of emotions. The OCC structure of emotions defines a hierarchical organization for emotion types. An emotion type represents a family of related emotions differing in terms of their intensity and manifestation. Thus, when an emotion type is referred, such as Fear, it does not specify the particular emotion associated with the word fear. Instead, it references the possible set of emotions resulting from appraising the prospect of a goal expected to fail, with varying degrees of intensity - concern, fright, petrified. Therefore, the attributes considered in the proposed model for the description of an emotion are:

- *Type:* The type of the emotion being experienced
- *Valence:* Denotes the basic types of emotional response. Positive or negative value of reaction
- *Target:* The name of the agent/object targeted by the emotion
- *Cause:* The event/action that caused the emotion
- *Intensity:* The intensity of the emotion
- *Time-stamp:* The moment in time when the emotion was created or updated

Every emotion has associated an Intensity attribute which is assigned with different values depending on the different situations that generated the particular emotion. Basically, it assesses how strong the emotion is. However, the intensity of an emotion does not remain constant during its life cycle in the system. Since the moment it is generated, the intensity of an emotion must be attenuated through time in order to reflect the dynamics of the emotional system itself. This characteristic reflects the notion that an emotion does not last forever and does not affect the evaluation of the subsequent emotional states in the same way. According to this concept, the model uses a decay function for emotions proposed by Picard [18], which characterizes intensity as a function of time. At any time (t), the value for the intensity of an emotion (em) is given by:

$$Intensity\,(em, t) = Intensity\,(em, t_0) \times e^{-b.t} \quad (1)$$

The constant b determines how fast the intensity of this particular emotion will decrease over time. This value can be controlled in order to reflect the short or long duration of the emotion types. The value *Intensity(em, t0)*, refers to the value of the intensity parameter of the emotion (em) when it was generated. When after some time t, the value of *Intensity(em,t)* reaches a defined threshold near zero, the emotion *(em)* must be removed from the system's repository, meaning that that specific emotion will no longer be part of the agent's emotional state.

In addition to emotions, the proposed model represents arousal and mood. Arousal represents the degree of excitement of the character. Aroused characters will feel more intense emotions. FearNot! only models psychological arousal, so whenever the character experiences a high intensity emotion (positive or negative), his arousal level will rise. Just like ordinary emotions, the arousal decays over time. Therefore, if nothing happens for a while, the character will "calm down". Mood represents an overall valence of the character's emotional state and is also used to influence the intensity of emotions. The idea, based on Picard, is that characters with a bad mood will tend to experience more negative emotions, and characters with a good mood will experience more positive emotions. Mood is represented as an internal variable that increases when positive emotions are created and decreases with negative emotions. This variable also decays over time until it reaches its neutral value (using a linear decay function).

4 Personality

In our model, the character's personality is also strongly based in OCC and is defined by: a set of goals; a set of emotional reaction rules; the character's action tendencies; emotional thresholds and decay rates for each of the 22 emotion types defined in OCC.

Our model uses two of OCC goal types, active-pursuit goals and interest goals. Active-pursuit goals are goals that the characters actively try to achieve, like going to a dentist appointment. Interest goals represent goals that a character has but does not pursue, as for instance wanting his favourite team to win a match, or avoiding getting hurt.

The emotional reaction rules assess how generic events are appraised and represent the character's standards and attitudes. Since the appraisal process is clearly subjective, these rules must be very dependent on personality. Action tendencies represent the character's impulsive actions which he performs without thinking (reactive actions). This labelling of reactive actions as action tendencies is due to Lazarus [13], which states that action tendencies are innate biological impulses, while coping is "a much more complex, deliberate and often planful psychological process". However, note that other psychological theorists may have distinct notions for the concept of action tendencies (ex: Frijda [9]). Specifying action tendencies for characters is very important to convey the viewer a

well defined personality. Loyall [14] pointed out that in order to achieve believability, characters must have very particular details of movements, mannerisms and reactions. As example, in FearNot!, when the victim is very sad it will tend to cry, while the bully will express his sadness in a completely different way.

OCC specifies for each emotion type an emotional threshold and decay rate. An emotional threshold specifies a character's resistance towards an emotion type, and the decay rate assess how fast does the emotion decay over time. When an event is appraised, the created emotions are not necessarily "felt" by the character. The appraisal process determines the potential of emotions. However such emotions are added to the character's emotional state only if their potential surpasses the defined threshold (the threshold is specific for each emotion). And even if they do overcome the threshold, the final emotion intensity is given by the difference between the threshold and the initial potential.

So, in addition to goals, standards and attitudes, these emotional thresholds and decay rates are used to complement a character's personality. For example, a peaceful character will have a high threshold and a strong decay for the emotion type of Anger, thus its anger emotions will be short and low. Thus, it is possible to have two characters with the same goals, standards and behaviours that react with different emotions to the same event (by having different thresholds). In order to model the decay rate, each emotion type has a different decay function (1), which differs in the constant value b. This value is given by the character's decay rate for each emotion.

5 Architecture

Figure 1 shows the complete model for our architecture. Taking into account the requirements depicted in the beginning of the paper, it presents two layers for the appraisal and coping processes. The reactive layer is responsible for the character's action tendencies, while the deliberative layer achieves the agent planful behaviour.

Action tendencies represent hardwire reactions to emotions and events that must be rapidly triggered and performed. Thus, the character must be able to react to an event and execute an action tendency almost immediately. Since the action tendencies depend on the character's emotional state, such assessment can only be made after the appraisal process. However, the cognitive appraisal depends on the agent's plans and can take some time: when an event is received, the continuous planner has to update all active plans (according to the event) even before the start of the generation of emotional reactions. For that reason, we applied the same two-level distinction to the appraisal process. Note that the planner will not extend (develop) any plan at this stage, it will just remove executed actions and update the plan probabilities accordingly.

While the deliberative level generates prospect-based emotions (hope, fear, satisfaction) based on the agent's plans and goals, the reactive level generates all other types of OCC emotions (fortune of others, well being, attribution, attraction) using a set of domain dependent emotional reaction rules as used by

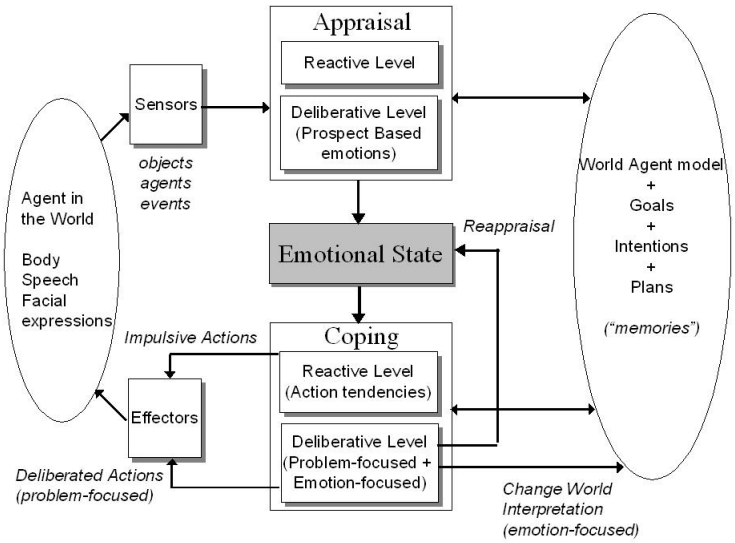

Fig. 1. Architecture Diagram

Martinho in S3A [16]. When an event is perceived, the reactive appraisal matches the event against the set of defined emotional rules, generating the corresponding emotions.

5.1 Cognitive Appraisal (Focus and Attention)

A continuous planner [20] that uses partial-ordered-plans builds up the core of the deliberative layer. However, the planner was extended to include probability information about actions and to perform emotion-focused coping strategies. The probability of a given action is biased by the character's personal interpretation and can be changed by emotion-focused strategies. More details about how the plans are represented, how a plan's probability of success is determined and how the planner works can be found in [5].

Each character has defined a set of active-pursuit goals that are triggered upon certain conditions. Thus, every time the agent receives a new perception from the environment, the deliberative layer checks all deactivated goals to determine if any of them has become active. If so, an intention to achieve the goal is added to the intention structure. The intention represents the agent's commitment to achieve the goal and stores all plans created for it. Initial hope and fear emotions based on the goal's importance are created in this process (and are stored inside the intention). After this initial process, the deliberative layer must choose between the existing intentions/goals to continue deliberation (and planning).

The OCC theory of emotions does not specify how do emotions affect reasoning/cognition and action selection. Thus, we had to look at the work of researchers that explored the influence of emotion in behaviour. For instance, according to Sloman [21], emotions are an efficient control mechanism used to detect situations or motives that need urgent response from the agent, and to trigger the appropriate redirection of processing resources. Applying the same idea in our architecture, we can use emotions to determine the most relevant intention: the ones generating the strongest emotions are the ones that require the most attention from the agent, and thus are the ones selected by the planner to continue deliberation.

After selecting the strongest intention, the best plan built so far is brought into consideration. This process is named focus and generates the following prospect based emotions:

- **Hope of success:** Hope to achieve the intention. The emotion intensity is determined from the goal's importance of success and the plan's probability of success.
- **Fear of failure:** Fear for not being able to achieve the intention. The emotion intensity is determined from the importance of failure and the plan's probability of failing.
- **Inter-goal fear:** Fear for not being able to preserve an interest goal. This emotion is generated if the plan contains any inter-goal threat.

In addition to active-pursuit goals, a character also has interest goals that specify protection constraints. These allow the modelling of conditions that the character wishes to protect/maintain. Whenever an action is added to a plan, a conflict between the action's effects and existing protected conditions may arise. This conflict is named an inter-goal threat. When the best plan is brought into focus, if it has any inter-goal threat, in addition to the normal emotions, it also generates a inter-goal fear emotion according to the respective interest goal that is being threatened. This emotion's intensity depends on the likelihood of the threat succeeding and on the interest goal's importance.

In the final phase of the deliberative appraisal, all active goals are checked to determine whether they succeed or fail. When such events occur or if the planner is unable to make a plan, more prospect based emotions will be generated, such as *Satisfaction*, *Disappointment*, *Relief* and *Fears-Confirmed*.

5.2 Coping

The coping strategies performed over the selected plan depends on the character's emotional state and personality. Inspired by CBI[15], the proposed model uses two types of coping: problem focused coping and emotional focused coping. The first one focuses on acting on the environment to cope with the situation, thus it consists on planning a set of actions that achieve the pretended final result and executing them. The second works by changing the agent's interpretation of circumstances (importance of goals, effect's probability), thus lowering strong negative emotions. When the planner analyses a plan, it applies every coping

Table 1. Applying coping strategies

Activation condition	Strategy	Effect
Plan probability very low	Acceptance	Drop the plan
Inter-goal threat detected, current goal's emotion stronger than interest goal's emotion	Acceptance or Wishful thinking	Accept the failure of the interest goal (ignore the threat) or lower the threat's probability
Inter-goal threat detected, interest goal stronger than goal's emotion	Acceptance	Drop the plan
Acceptance strategy applied	Mental Disengagement	Lower the goal's importance
Causal Conflict detected	Planning or Wishful thinking	Use promotion, demotion, or lower the conflict probability
Open Precondition	Planning	Add a step that achieves the precondition
Consistent plan without open preconditions	Execution	Execute an action

strategy that satisfies its conditions (with a specific order). Table 1 presents the several coping strategies.

Acceptance is the recognition that something is not possible to achieve or protect/maintain. If the selected plan's probability is lower than a given threshold, the character thinks that it's not worth the time to try to improve the plan, since adding more actions will not increase its probability of success, and drops the plan.

Whenever an acceptance strategy is applied, mental disengagement is also applied. Mental disengagement works by reducing the goal's importance. Since acceptance will frequently lead to goal failure, lowering the goal's importance reduces the intensity of the negative emotions triggered when the goal fails. This does not mean that to fail is ok, in fact the character will feel distressed about failing. Mental disengagement just slightly mitigates his distress.

If the planner detects an inter-goal threat in the plan, it can use additional coping strategies. If the threatened condition generates emotions stronger than the goal's emotions, the current plan is dropped. In the opposite situation, the character can either accept the interest goal's failure (by removing the protected condition) or use wishful thinking to cope with the fear emotion. Wishful thinking works by denying the reality of an event or by thinking that something bad will not happen. This strategy lowers the threat's probability by lowering the probability of the effect that threatens the condition.

Finally, when the planner achieves a consistent plan with no open preconditions it has reached a solution. This solution that corresponds to a partial ordered plan is then executed by repeatedly choosing and performing any of the next possible actions.

It is important to point out that since part of the coping strategies are triggered by emotions, the emotional state and personality influence the strategies applied and hence the overall reasoning performed by the characters. For instance, a fearful character has much more chances to drop an active pursuit goal if it presents threats to other goals.

6 Illustrative Example

This section presents an illustrative example that helps the reader understand how the mechanisms described in the previous section achieve a believable behaviour. In FearNot!'s first episode, John, the victim is studying in the classroom when Luke (the bully) enters. Luke does not like John and so when sees John he starts insulting him. As a result, John has an active pursuit goal of fighting back that is triggered when he is insulted by other characters. So, he tries to build a plan in order to fight back. However all the actions that John considers to fight back have some likelihood of getting hit back. So, when such plans are selected, a threat to John's interest goal of not getting hurt is detected and thus John feels frightened. Due to the victim's fearful nature, his inter-goal 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 John is thinking, Luke notices the book over the table and realizes a bullying opportunity. So he makes a plan to push John's books to the floor. Luke moves towards the books and pushes them away. This event matches an emotional reaction that generates a gloating emotion, which triggers an action tendency. Luke performs a tease speech act that corresponds to saying something along the lines: "Come and get them you Muppet!" When the victim realizes that the books are on the floor he activates the goal of picking them, and thus walks towards them and picks them up. As if it was not enough, when the bully sees John picking up the books he decides to push him. So Luke goes behind John and pushes him (see Fig. 2-a).

When 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 to gloat about John by saying something along the lines of "What a wimp, I've hardly touched you". When John cries, Luke finds it very blameworthy and thus threatens him to stop crying and to not tell anyone.

At this point, the episode ends and the application changes to the interaction mode (Fig. 2-b) where the child user talks with the victim (John) and advises him on what to do next. The agent perceives three suggestions: fight back, ignore the bully or tell someone. These strategies correspond to goals that the character already has but is unable to achieve because they usually threaten other interest goals. Suppose that the user types something containing "hit" or "kick" or "punch", the language system recognizes such utterances as fightback suggestions. The agent remembers his goal's results and knows that he has not

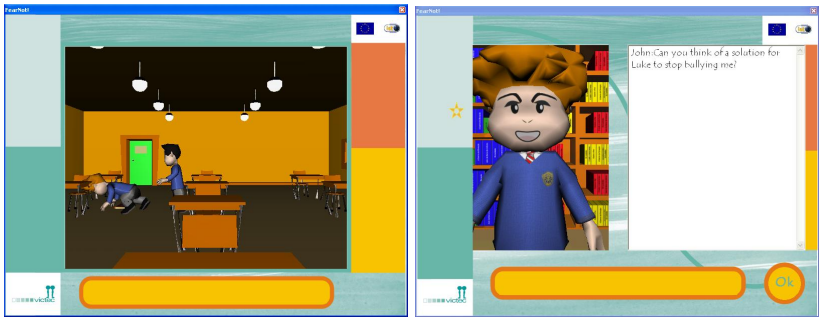

(a) Paul pushes John (b) The user interacts with John

Fig. 2. FearNot! application

actually tried to fight back before, so he accepts the suggestion and increases the fight back goal's importance.

Therefore, on the next episode, when John is insulted once more, he will activate his goal of fighting back. However, since the goal's importance is much higher now, the emotions created by the goal are stronger than the ones created by the threat of getting hurt. So, instead of giving up fighting, John decides to confront Luke.

In the victim's perspective, both user and character's actions are perceived equally. Therefore, the victim can experience emotional reactions triggered by the user's interaction. For instance, the victim feels satisfied whenever he receives a suggestion from the user.

7 Evaluation

In order to obtain some results concerning the believability of the characters created and the emotional reactions be the children, we performed a preliminary evaluation with the goal of determine whether the architecture is capable of achieving similar results as the scripted version concerning believability and empathy felt by the children.

The scripted version used the same characters as the emergent one, but all the behaviours were predefined for each episode. After each episode, the victim appears to the user like in the emergent version, however the interaction is different: the child just has to select between different strategies (by pressing a button).

Experiments made with the scripted version in Portuguese, English and German schools showed that the children found the characters believable and felt empathy for them [4]. Although we have not yet made a similar cross-cultural evaluation with the emergent version, the preliminary evaluation gave us some insights on how children react to the autonomous characters. For the experiment,

we chose eleven children aged between eight and eleven from third and fourth year, male and female from a Portuguese school.

The results obtained were similar to the scripted version. As expected, Luke (the bully) is the most disliked character, while John (the victim) is the favourite one. This means that children really create an empathic bond with the victim character just like in the scripted version. In order to assess if children felt any emotional reactions to the situations created with the characters we tested two reactions: (1) did they feel sorry for any character? And (2) did they feel angry at any character? Figure 3 shows that children did feel sorry at John (the victim), and reported feeling angry at the bully character, as expected. The additional character, Paul (a friend of John), appears in one of the episodes when John asks for help. These results show that the architecture's emergent behaviour can also elicit emotional reactions from the users.

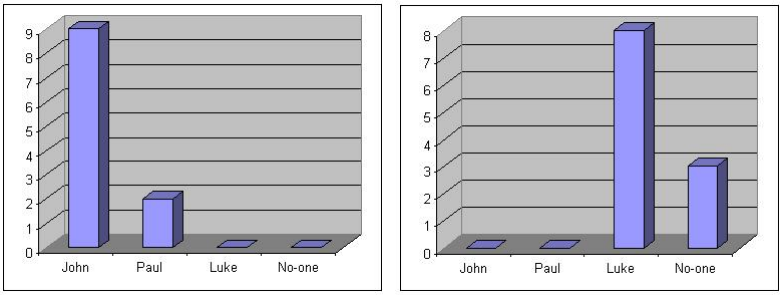

(a) Whom did you feel sorry for? (b) Whom did you feel angry at?

Fig. 3. Emotional reactions from the users

In addition, we asked children if they felt in control of the story and if they liked to interact with the victim. These questions were evaluated using a Likert scale (1-5; 1 - Very good; 5 - Very Bad). Table 2 shows the questions and the results obtained with both versions.

Table 2. Scripted Vs. Emergent version

	Scripted	**Emergent**
Did the conversations seem real? (yes-1;no-5)	2.4	1.9
Where 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

We believe that the conversation with the victim, and the fact that the victim follows the children suggestions, makes children find the overall dialogue and narrative seem more real and interesting. For example, when the victim accepts the fight back strategy, it seems more real to see him threatening the bully on the next episode than to behave like in the first episode.

8 Final Remarks

The results attained, although very limited, do however suggest that the use of autonomous synthetic characters can lead to believable situations that do evoke empathy in users, and that like in traditional animation, emotions and personality are key to achieve this goal. Furthermore, the results also suggest that the use of the characters is an advantage as far as building a conversational interaction with the users. The fact that the children help the victim by giving him suggestions creates a stronger bond between the character and them. For instance, if the victim succeeds in coping with the situation both the victim and the child will feel satisfied. The child feels satisfied not only as an empathic response but also because she took an active role in the situation's success.

In the future, we expect to perform more tests, in particular with more parameterizations, and for relational bullying scenarios. Indeed, the use of the agent architecture poses no problems to building and configuring those new episodes.

References

1. Bates, J.: The Nature of Character in Interactive Worlds and The Oz Project. Technical Report CMU-CS-92-200, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1992.
2. Bates, J.: Virtual Reality, Art, and Entertainment. Presence: The Journal of Teleoperators and Virtual Environments, 1(1), 133-138, 1992.
3. Bates, J.: The Role of Emotion in Believable Agents. In Communications of the ACM, Volume 37, Issue 7, 122-125, July 1994.
4. Zoll, C., Enz, S., Schaub, H., Woods, S., Hall, L.: Deliverable 7.2.1 Evaluation of Demonstrator in Schools, VICTEC Project, 2005
5. Dias, J.: FearNot!: Creating Emotional Autonomous Synthetic Characters for Empathic Interactions. Universidade Técnica de Lisboa, Instituto Superior Técnico, Lisboa, MSc Thesis, 2005.
6. Dautanhahn, K., Woods S.: Possible connections between bullying behaviour, emphaty and imitation. Thecnical report, 2003
7. Eisenberg, N., Strayer J.: Critical issues in the study of empathy. In N. Eisenberg and J. Strayer, editors, Empathy and its Development. Cambridge University Press, 1987.
8. Elliott, C.: The Affective Reasoner: A process model of emotions in a multi-agent system. Northwestern University, PhD Thesis, Illinois, 92.
9. Frijda N.: The Emotions. Cambridge University Press, UK, 1986.
10. Gadanho, S.: Learning Behavior-Selection by Emotions and Cognition in Multi-Goal Robot Task. In Journal of Machine Learning Research 4, pages 385-412, 2003
11. Gratch, J.: Émile: Marshalling Passions in Training and Education. In 4th International Conference on Autonomous Agents, ACM Press, June 2000.
12. Gratch, J., Marsella, S.: A Domain-independent Framework for Modeling Emotion. Appears in Journal of Cognitive Systems Research, Volume 5, Issue 4, pages 269-306, 2004.
13. Lazarus, R.: Emotion and Adaptation. Oxford University Press, New York, 1991.
14. Loyall, B.: Believable Agents: Building Interactive Personalities. Technical Report CMU-CS-97-123, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1997.
15. Marsella, S., Johnson, L, LaBore, C.: Interactive Pedagogical Drama. In proceedings of the Fourth International Conference on Autonomous Agents, pp. 301-308, ACM Press, 2000

16. Martinho C.: Emotions in Motion: short time development of believable pathematic agents in intelligent virtual environments. Universidade Técnica de Lisboa, Instituto Superior Técnico, Lisboa, MSc Thesis, 1999.
17. Ortony, A., Clore, G., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press, UK, 1988.
18. Picard, R.: Affective computing. MIT Press, Cambridge, 1997.
19. Reilly, S.: Believable Social and Emotional Agents. Ph.D. Thesis. Technical Report CMU-CS-96-138, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1996.
20. Russel, S., Norvig, P.: Artificial Intelligence: A Modern Approach. 2 Edition Englewood-Cliff, N.J.: Prentice-Hall, 2002.
21. Sloman, A.: Damasio, descartes, alarms, and metamanagement, In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC '98), pages 2652-7, San Diego, CA, USA, 1998
22. Staller, A., Petta, P.: Towards a Tractable Appraisal-Based Architecture for Situated Cognizers. In Grounding Emotions in Adaptive Systems, Workshop Notes, Fifth International Conference of the Society for Adaptative Behaviour (SAB98), Zurich, 1998.
23. Thomas, F., Johnston, O.: Disney Animation: The Illusion of Life. Abbeville Press, New York (1981)