An affectively driven planner for synthetic characters

Ruth Aylett[1], Joao Dias[2], Ana Paiva[2]

[1]School of Maths and Computer Science
Heriot-Watt University
Edinburgh E14 4AS UK
Ruth@Macs.hw.ac.uk

[2]INESC-ID
IST, Taguspark
Porto Salvo, Portugal
joao.assis@tagis.utl.ist.pt Ana.paiva@inesc-id.pt

Abstract

This paper discusses the requirements of planning for believable synthetic characters and examines the relationship between appraisal and planning as components of an affective agent architecture. It discusses an implementation in the synthetic characters of the FearNot! anti-bullying education demonstrator and how far this provides an adequate mechanism for believable behaviour.

Introduction

Synthetic characters, or intelligent virtual agents, (IVAs) are graphically embodied agents that are able to act autonomously in their virtual world. They therefore have similarities to robots in that they must implement individual sense-reflect-act cycles and interact (pseudo-) physically with their environment. The real-time constraints are if anything more demanding than in robotics (many robots move extremely slowly) but the need to interleave planning and execution so that there is always some sensible action to carry out is very much the same. However there are also some important differences.

Some of these are obvious – for example sensing and actuation are far less complex since a virtual world has only as much physics as its designer chooses to build in. Thus in sensing, the processing overload and inherent difficulty of deriving symbolic information from the world is very much lower – sometimes non-existent - and in actuation, IVAs do not run out of battery power or topple over when walking. A consequence is that they are able to carry out a more substantial repertoire of more complex actions than is feasible for robots, making high-level planning (as distinct from motion-planning) a more useful component of their overall architecture.

A less obvious difference however concerns the role of IVAs and the metrics for success. While robots were until recently almost entirely task-oriented, with execution robustness and efficiency the key requirement, IVAs have from the beginning been measured by the attribute of *believability*. This term, generally derived from the seminal work of the Oz project [6] is intuitively clear but hard to formally define. It can be taken as the extent to which IVAs allow human users interacting with or observing them to suspend their disbelief and treat them as convincingly lifelike with an inner life of their own. Note that believability is not at all the same thing as naturalism: Mickey Mouse is a believable character but is not at all like a real mouse. IVAs are often not task-oriented in the robot sense, but can be seen as social actors [11] that are expected to 'live' believably in their own virtual world and act as socially competent interaction partners with human users.

An important aspect of maintaining believability is communicating the internal state of the character to a human user. This supports the continuing human process of inferring the intentions of an IVA – its motives and goals - and can help to produce the impression of coherent action which is required for the user to feel that in some sense they 'understand' what an IVA is doing. Of course witnessing the execution of the planned actions of the IVA is one means of doing this, but an important aspect of embodiment is its use as an extra communication mechanism for the internal state of the IVA through *expressive behaviour*, complementing its explicit communication mechanisms of action execution and natural language [8]. Extra mechanisms, either reactive or involving planning, are needed to handle this.

An important aspect of expressive behaviour is of course emotional expressiveness. A user will try to interpret this as revealing something about the emotional state of the IVA which they can relate both to the IVAs motives and goals and to their own affective state. Indeed, if IVA expressive behaviour lacks the appropriate affective dimension, its believability is severely undermined, underlining the important role that emotion must play in IVAs. While emotional overtones can be mechanically scripted onto IVA animations, the most straightforward way of generating appropriate expressive behaviour lies in incorporating an affective model [3].

Alongside this external behavioural requirement, work with human subjects has shown that emotion is by no means as disjunct from cognition as was assumed in Descartian dualism [10] but can be seen as one aspect of an integrated system, with particular relevance to attentional focus, goal selection and plan evaluation. Thus a further requirement for a planning system in an IVA is that it be integrated with the affective model.

Overall we see that the requirements for planning in IVAs are in some ways very unlike those in conventional robotics, though very similar issues are in fact being raised in the new field of Human-Robot Interaction (HRI) and social robotics [12].

VICTEC and Interactive narrative

As a concrete example of implementing planning in IVAs, we discuss work carried out in the European Framework V project VICTEC (Virtual Information and Communication Technology with Empathic Characters) [24]. This project sought to produce a system to help with anti-bullying education – and by extension, other areas of Personal and Social Education (PSE) – by building empathy between a child user and a synthetic character in a virtual drama. These dramas involved short episodes in which one character was bullied by one or more other characters, with the child user acting as an 'invisible friend' and trying to help a victim [7].

Believability was a very significant issue since a premise of the project was that the creation of empathy requires the user to feel the characters have an independent life, that the events 'really' happen to them, and really affect them emotionally, in a way quite different from the indestructible or infinitely regenerating protagonists of most computer games.

This suggested a need for unique narratives, that is, narratives with different characters and events for different users, rather than scripted stories that repeat identically. If stories are literally repeatable, then one loses the sense that the characters have any control over their virtual lives, while if the same characters repeat different narratives (possibly due to user intervention) then the coherence of the character with which the user is to feel empathy is lost. The mechanism chosen for continuing but different narratives was *emergent narrative* [1,2], that is, narrative generated by interaction between characters in the style of improvisational drama, rather than the authored narratives in more widespread use.

Bullying is naturally episodic and while each time is different in some sense each time is also the same, thus not offering too great a challenge for an initial implementation of an emergent narrative approach. The requirement that the child influence the character underlines the need for such an approach, since branching on every possible suggestion over a number of episodes would otherwise produce a combinatorial explosion, while the child soon notices if a scripted agent in fact takes no notice of their advice.

If interesting stories are to be generated by interaction between characters in this way, then the agent architecture of the IVAs must meet the challenge of producing sequences of actions with accompanying expressive behaviour that meet the dramatic situations in which they are located. The bully must be motivated to carry out aggressive behaviour sequences towards the victim, and the victim must respond in a believable way but with some influence from the advice the child acting as their friend has given them. It is also very clear that the emotional reactions of the IVAs to each other provide the motivational structure for planning as well as the basis for unplanned reactive behaviour such as crying.

The agent framework developed for the project allows construction of IVAs that express and react to emotions in a natural and meaningful way,. It has been designed so that it does not only apply to the specific context of school bullying, but can be used in the more general realization of emergent dramas.

The FearNot! agent framework

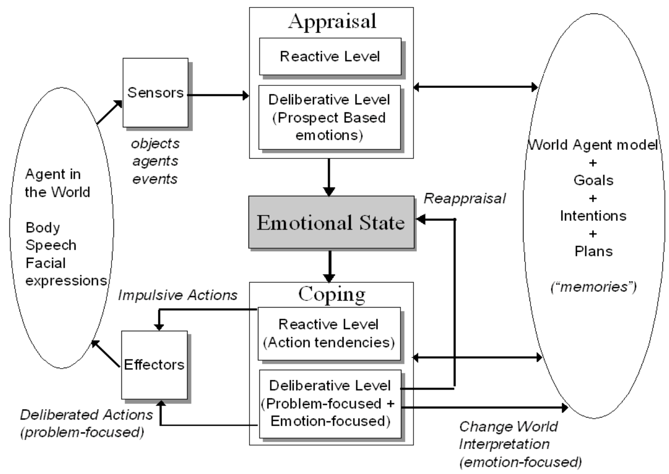

Figure 1: Affective agent architecture

The agent architecture used in the FearNot! demonstrator (Fun with Empathic Agents to Reach Novel Outcomes in Teaching) is shown in **Figure 1**. Emotional status affects an agent's drives, motivations, priorities and relationships. Key concepts in this architecture are those of appraisal [14,17] and coping behaviour [17], both taken from cognitive appraisal theories of emotion. Appraisal is the perceptual process through which objects, other characters and events are related to the needs and goals of an IVA with an emotional response generated as a result. Coping behaviour is the internal emotional adjustment made or external actions taken in order to deal with negative emotions; in the latter case it invokes planning.

Inspired by hybrid agent architectures [19], FearNot! provides two distinct levels in both appraisal and coping mechanisms. The reactive level provides a fast mechanism to appraise and react to a given event, while the deliberative level takes longer to react but allows a more sequentially complex and rich behaviour.

Appraisal and planning

The emotional framework implemented in FearNot! is a subset of the one proposed by Ortony, Clore and Collins (OCC) [20]. The OCC model is an approach based on a valenced (good or bad) reaction to an event and the structure of emotions it defines can be seen as a hierarchical taxonomy which organises 22 emotion types. It first divides emotions into groups as seen in Table 1:

To define..	Is to state…
Attraction reactions	Which entities the character likes or hates
Event reactions	What importance specific events have for a character
Prospect-based reactions	What importance future events have for a character
Attribution reactions	What standards of behaviour are important to a character: praiseworthiness

Table 1: OCC emotion groups

By decomposing this hierarchy further, one can define a set of appraisal rules relating events, other agents and time to emotional reactions, as for example in Table 2:

Sub-group	Appraised as	Sample rules
Well-being	event	Event congruent with goals – joy Event not congruent with goals - distress
Fortunes-of-others	event affecting another	Happy-for pleased about an event desirable for another Resentment displeased about an event desirable for another
Prospect-based	prospective event	Hope pleased about a prospective desired event Fear displeased about a prospective undesired event

Table 2: Example OCC appraisal rules

This approach has been widely implemented, partly because it is both comprehensive and computationally straightforward, though it is worth pointing out that it was designed as an account of how people reason about the emotions of others rather than how their own emotions are generated. In this work, it was used to represent emotions as shown below in Table 3.The type attribute refers to the generic type of the emotion experienced: fear, anger, joy etc. Each emotion type can be realized in a variety of related forms with varying degrees of intensity (i.e. emotion type Fear can generate an emotion range from 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 emotional impact and potential response to the stimulus.

Attribute	Description
Type	The type of the emotion being experienced
Valence	Denotes the basic types of emotional response (positive or negative)
Target	The name of the agent/object towards the emotion is directed
Cause	The event/action that caused the emotion
Intensity	The intensity of the emotion. A logarithmic scale between 0-10
Time-stamp T_0	The moment in time when the emotion was created

Table 3: VICTEC agents' emotion attributes

To deal with the dynamic aspect of emotion, the intensity of an emotion decays from the moment it is generated onwards using Picard's [20] decay function defining intensity as a function of time. At any time (t), the value for the intensity of an emotion is (em) is given by the formula:

$$\mathrm{Intensity}_{em,t} = \mathrm{Intensity}_{em,t0} \times e^{-dt}$$

Here d (decay) determines how fast the intensity of this particular emotion will decrease over time. The value Intensity(em, t_0), refers to the value of the intensity parameter of the emotion (em) when it was created.

Note that the decay factor is not only different for different emotions, but can also be differently set for different characters. If this is then combined with a threshold value below which a character does not respond to the emotion even if an appraisal generates it, then the effect is to produce different personality types without having to explicitly model personality.

The appraisal mechanism

As shown in *Figure 1*, the appraisal mechanism is composed of two distinct layers. The reactive layer appraisal is handled by a set of emotional reaction rules, based on Elliot's Construal Theory [13]. A reaction rule consists of an event that triggers the rule and values for OCC appraisal variables affected by the event (desirability, desirability-for-other, praiseworthiness etc).

The deliberative layer is responsible for appraising events according to the character's goals, thus generating prospect-based emotions like hope and fear. Because these emotions specifically relate to future events – either to those congruent with the IVA's goals (hope) or threatening those goals (fear), they offer a specific interface between the affective system and the planning component of coping behaviour.

As Gratch [16] first proposed in the system Émile, the state of plans in memory can be used internally to generate prospect based emotions. When a given plan is considered by the planner it will generate emotion potentials according to the formulae:

*HopePotential = P(Plan) * ImportanceOfSucess*
*FearPotential =(1 - P(Plan)) * ImportanceOfFailure*

Here P represents the probability of success of a plan, and ImportanceOfSucess/Failure relates to the importance of the goals for which the plan has been built.

Although emotion potential is strongly related to classic decision theory information, there is an important difference. The final intensity of emotions is biased by personality, supporting a greater differentiation of behaviour between different characters. The stronger of the two emotions hope and fear is called the *dominant* emotion and is used to guide planning as described later. A fearful character has a low threshold and experiences *Fear* more easily, making this the dominant emotion more often. It therefore considers goals that seem unachievable earlier, and gives up goals that threaten other interest goals much more easily. On the other hand, a hopeful character is usually driven by *Hope*, producing in general a more optimistic and bold behaviour.

Notice that although it seems contradictory, a character may experience both hope and fear at the same time (driven by the same goal). The OCC Theory in fact supports the existence of simultaneous contradictory emotions. For example consider someone whose uncle has recently died but left him a huge legacy; one could understand the co-existence of contradictory *Joy* and *Distress* emotions. In the same way, humans may experience hope and fear at the same time.

FearNot! includes two of the goal types defined in the OCC taxonomy; *active-pursuit* and *interest goals. Active-pursuit* goals are goals that the character actively plans to achieve (i.e. going to a dental appointment) while *interest goals* represent goals that the character has but does not plan specific sequences of actions to achieve (i.e. avoiding getting hurt). The OCC *replenishment goals* are not used since they could be considered as *active-pursuit goals* with cyclic activation and deactivation.

An *active-pursuit goal* is defined in the VICTEC agent architecture by the following attributes;

Id *(goal identifier),*
Type (active-pursuit),
Pre-conditions,
Success-Conditions and Failure-Conditions.

Unlike the *active-pursuit goal*, the *interest goal* does not have any pre-conditions, success or failure conditions since it does not become active or inactive. However the *interest goal* possesses one extra parameter; a protection-constraint. This supports modelling of conditions that the character wishes to maintain, so the planner will try to prevent actions that threaten such conditions.

The core of the deliberative layer is a partial-order continuous planner [22] that was extended to include emotion-focused coping strategies. The deliberative layer has to monitor all events in order to detect when an action is accomplished or fails. It can also handle unexpected events that affect future plans and it can handle serendipity: suppose that the planner has finished building a plan to achieve a goal - if some other agent comes in and achieves some precondition for us, the planner will detect that the condition holds true in the start step and will remove the action used to achieve such a precondition.

Thus the first step in the deliberative appraisal is to update all plans accordingly to the event being appraised and also to update the probability of action effects succeeding. If an action was successfully executed but an expected effect did not occur, the planner updates effect probability accordingly. This process will change the agent's internal plans (and plan probabilities) leading to different emotional appraisals of Hope and Fear.

In addition, when an event is appraised, the deliberative level checks if any goal has become active, and if so, an intention to achieve the goal's success conditions is created generating initial hope and fear emotions.

The action selection and coping mechanism

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

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

The deliberative coping process is more complex and is deeply connected to the deliberative appraisal process just discussed. More than one goal can be active at the same time, so the first stage of the deliberative reasoning process is to determine which goal to attend to. According to Sloman [23], 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. The idea is that the intentions 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.

Once the most relevant intention is selected, it is necessary to choose the most appropriate plan from those already existing to execute or to continue planning. The best plan is given by the following heuristic function:

$$h(plan) = (1 + numberOfSteps + numberOfOpenPreconditions + numberOfInterThreats * 2) / P(plan)$$

Using the number of steps gives a better value to plans that achieve the same conditions but use fewer steps (this is likely to be a better plan). Between plans that have the same number of steps and open preconditions, the ones with higher probability have a lower value of h and are therefore better. In addition to the number of steps and the number of open preconditions, this function has one additional parameter: the number of inter-goal threats in the specific plan. When a plan has an action that violates a protected condition specified in an interest goal, that action is considered a inter-goal threat. With this heuristic function, plans that have such threats are less often chosen.

At this point, the best plan is brought into focus for reasoning, as if in the forefront of the agent mind. When a plan is brought into focus,- and only in this situation - it generates/updates the corresponding emotions. So, if a very unlikely plan is created during the planning process and added to the intention structure, it does not generate a strong fear emotion unless it is forefronted in this way: one can interpret this as the character not worrying because it has another, better plan in mind.

When a plan is brought into focus, it generates the following prospect based emotions:

Hope: Hope that it may be able to achieve the intention.

Fear: Fear that it may not be able to achieve the intention.

InterGoal-Fear: Fear that it may not be able to preserve an interest goal. This emotion is generated if the plan contains any inter-goal threat.

As mentioned previously, an interest goal may 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 *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.

The planner is only allowed to remove one flaw or start to execute one action in each cycle of coping, so that an IVA does not 'freeze' in prolonged thought. Thus building up a plan takes several coping cycles, which means that several appraisals will be made for the same plan. In fact, an initial appraisal may drastically change from an initial strong hope to a strong fear emotion as the character realizes that no feasible plan exists. This type of appraisal is called *Reappraisal* since it is not based on a external event or stimuli, but driven by the IVA's internal processing.

The selected plan is then analysed by the planner, which uses the problem-focused and emotion-focused strategies discussed below to process and generate new plans. One or both may be applied in each coping cycle if appropriate. Unlike classic planners, the FearNot! planner can also use emotion-focused strategies to drop an unlikely plan, to improve a plan or to resolve a flaw. The resulting plan is stored with the intention and can be pursued later on.

Plan representation

Plan operators are a slight modification of STRIPS operators, associating probability values with the effects. A plan operator consists of the action with pre-conditions and effects attributes, where the pre-conditions list contains a set of conditions that must be verified in order to execute the action and the effects a list of conditions that will hold when the action ends.

Figure 2 shows the operators for two actions: the **walk-to** action and the **ask-question** action. In the first case, the agent must be standing in order to walk, and the action effect (with 100% probability) is that the character will be positioned in the target interaction spot, which is defined in the visualisation of the world. The probability is of course domain-dependent: in this world characters are not physically prevented from walking to a particular position. The second operator is an example of a language action. In this case the probability has been set at 0.5 because it can receive a positive or a negative answer.

As well as a set of operators, a plan includes ordering constraints, causal links, binding constraints and open pre-conditions in the standard way.

Emotion Focused Coping

Marsella and Gratch [15, 18] introduced the use of emotion-focused coping in planning processes. This is an internal process that works by changing the agent's interpretation of circumstances, thus lowering strong

```
<DomainActions>

  <Action name="walk-to([target],[interaction-spot])">
    <PreConditions>
      <Property name="?[SELF](pose)" operator="="
value="standing" />
    </PreConditions>
    <Effects>
      <Effect probability="1.0">
        <Property name="?IS([interaction-spot],[SELF],[target]"
value="True" />
      </Effect>
     </Effects>
  </Action>

  <Action name="AskQuestion([agent],helpprotection)">
    <PreConditions>
      <Property name="[agent]" operator="!=" value="[SELF]" />
    </PreConditions>
    <Effects>
         <Effect probability="0.5">
        <Property
name="?EVENT([agent],SpeechAct,[SELF],helpprotection,po
sitiveanswer"
         operator="=" value="True" />
      </Effect>
      <Effect probability="0.5">
        <Property
name="?EVENT([agent],SpeechAct,[SELF],helpprotection,ne
gativeanswer)"
         operator="=" value="True" />
      </Effect>
    </Effects>
  </Action>
</DomainActions>
```

Figure 2: Two examples of operators in FearNot!

negative emotions. It is often used by people, especially when problem-focused coping (which corresponds to acting on the environment) has a low chance of success – for example 'I never really wanted it anyway', 'there's no point in setting myself up to fail'.

The FearNot! deliberative planner uses the emotion focused strategies of *acceptance*, *denial* and *mental disengagement*.

Acceptance is the recognition that something cannot be achieved, and accepting failure. When a plan has a very low probability of success, adding new actions to it will only make it worse, so that the planner will accept that the plan will fail and will not try to improve it. If an acceptance strategy is being applied it means that the planner is considering plans with a low level of probability and there is a good chance that the goal will also eventually fail. For these reasons, *mental disengagement* is also used whenever acceptance is applied, and this works by lowering the goal's importance (so that the character will not feel as disappointed when the goal fails).

The most important role of acceptance is when a plan step threatens another goal (say an interest goal's protected condition). If the active pursuit goal generates stronger emotions than the interest goal, the plan is maintained and the protected condition failure is accepted. Otherwise, the plan will be dropped, again applying acceptance.

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

Since the agents in VICTEC are emotionally driven, any significant interaction with a child user or another agent will result in the alteration of the agent's emotional state. Since the agent makes decisions based on that emotional state, this potentially affects its perception of actions and alters the probability of plan success and the resulting feelings of hope and fear. This, in turn, influences the actions selected for execution by the agent and allows for the unfolding of narratives different in form and content (i.e. according to their context) without the need for scripting them.

An example

In this section we give an example of a run of the implemented system.

In the first episode, John, the victim, is located in the classroom studying and Luke enters (Figure 3). Luke does not like John and so when he sees John he starts insulting him (reactive action tendency). Note that the same mechanism that plans and selects physical actions also plans and selects language actions – each language action is then turned into an utterance by a template-based language system [4].

Figure 3: Luke enters the classroom

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

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

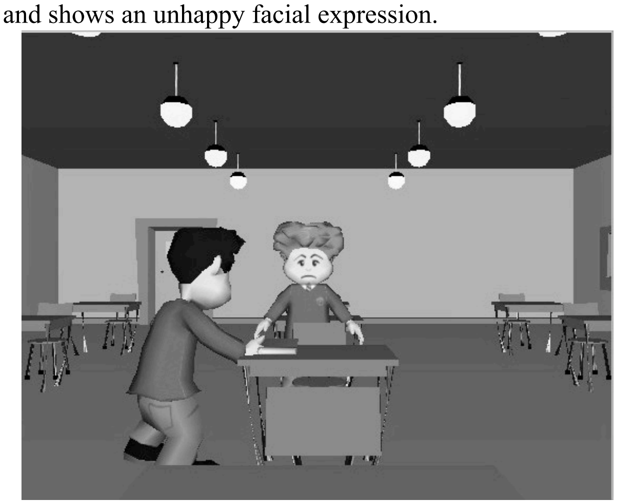

Figure 4: Luke pushes John's book off the desk

Figure 5: Luke pushes John: he falls

Luke moves towards the books and pushes them away (Figure 4). This event matches an emotional reaction generating the emotion *gloat*, which triggers an action tendency. Luke performs a *tease* language action that corresponds to saying something like: “Come and get them you Muppet!” When the victim realizes that the books are on the floor he activates the goal of picking them, and thus walks towards them and picks them up. When the bully sees John picking up the books he decides to push him. Once more this is achieved by an active pursuit goal that becomes active in that situation. So Luke goes behind John and pushes him.

The result of pushing John is uncertain: in the real world it is decided by physics, and in the virtual world by a probability set in the 3D visualization. Thus sometimes a character may fall (Figure 5), and sometimes not. If John falls, he appraises this event as very undesirable and activates an action tendency to start crying. At the same time, Luke appraises the same event as very desirable and starts gloating at John’s situation by saying something in the lines of “What a wimp, I've hardly touched you”. When John cries, Luke finds it very blameworthy and thus threatens him to stop crying and to not tell anyone. If John does not fall, Luke will not mock him. Instead, the victim may feel angry and asks Luke why is he always picking on him. Luke responds negatively to the question by insulting John even more.

Figure 6 shows a snapshot of the interaction mode in which the child user talks with the character victim and advises him/her on what to do next. The user types whatever they want in the lower text box and by pressing the OK button the written utterance is sent to the agent. The agent receives the utterance and converts it to a language action using the template-based language system, already mentioned, in reverse. When the interaction mode is first displayed, John arrives in the library crying, but he realizes that the user has entered the set as for any ordinary character (in fact the agent victim does not distinguishes the user from other synthetic agents) and activates the goal of asking for help which makes him perform an *askforhelp* speech act.

If the user then suggests fighting back, this has the effect of raising the importance of the goal, so that the next time John meets Luke the fear generated by the possibility of getting hurt is not strong enough to make him give up the goal. Thus user interaction changes the behaviour of the victim by indirect influence rather than because the victim does exactly what he is told. However if John tries pushing Luke and it does not succeed, then he will not accept a further suggestion to hit back since the experience of being hurt as a result again alters his emotional state, this time in the direction of greater fearfulness.

And does it work?

FearNot! was initially evaluated extensively in a scripted version in order to test the premise that child users would empathise with the problems of the victim character, the basis for meeting the pedagogical objectives . Space does not allow a detailed discussion of the results, which can be found in [26]. However, in summary, testing with 420 children in June 2004 demonstrated that in spite of reservations about the quality of graphics and animation, the use of scenarios captured from children themselves produced the high level of engagement and empathy that had been aimed for. Interesting gender differences were observed, with girls more able to empathise with male and female victims, while boys empathised much more with male than female victims.

Given that a scripted version – which did not therefore need the planning system described – was able to have this effect, an obvious question is, why have a planner at all? The answer lies in the restricted interaction style and narrative scope enforced by scripting. These stand up for a short once-only interaction, such as the 30-40 mins each child user experienced during the evaluation, but in a live school setting in which FearNot! would be used over a period of weeks in conjunction with other curriculum elements, we believe they would become a serious weakness.

The scripted system did not use the free-entry language interface of Figure 6, since in the nature of scripting, only one or at most two pieces of advice would actually have any effect on the trajectory of the story. Instead a randomly ordered menu allowed one of six pieces of advice to be selected, with free text added to explain why the choice had been made. Three episodes were in fact run, and only the third of these could change: if the child selected telling a friend, parent or teacher as one of their pieces of advice then the third episode showed a successful ending, and otherwise an unsuccessful ending. The rigidity of scripting thus meant that the victim character at times evidently did not act on the specific advice given, and this was noted by a number of children both in questionnaires and post-evaluation focus group discussion.

Given the combinatorial effect of different coping response choices as advice over multiple episodes, it is our view that scripting, or the use of branching narrative, cannot be successful over longer interactions. This resulted in the development of the unscripted system discussed in this paper, in which the planner is what allows characters to interact so as to create an emergent narrative. A small-scale evaluation of this work has so far been completed with 11 children, and this has shown an increase in the ratings given for character responsiveness, from the question 'Did the victim respond to your advice?" [27]. A longer evaluation in the school environment will take place during 2006-7.

Related work

As indicated above, this work draws specifically on that of Marsella and Gratch [15,16,18] in integrating emotion with planning in an overall affective agent architecture. This in turn develops the reactive architecture of Elliot [13] by adding predictive planning, allowing sequences of actions to be produced by agents rather than only one-shot reactions. This work goes further than that cited however by supporting interactions between agents with no explicit branching choices by the user determining the dramatic direction. Thus a greater onus is placed upon the planning mechanism.

As in Cavazza et al [9], interaction between the plans of agents can produce different dramatic outcomes, but unlike that work, FearNot! plans are generated at run-time rather than being explicitly represented as pre-authored hierarchical and-or trees. Cavazza et al allow intervention by the user at run-time to create some of the variability required to produce different stories: in FearNot! the child user is deliberately not given any direct means to intervene in the dramatic episodes, firstly, because it is not meant to be a 'bash the bully' game, and secondly, because seeing how the advice given influences the victim's behaviour is essential to promoting reflection and thus reaching the pedagogical outcomes. However an equivalent uncertainty is created both by the indeterminate outcome of physical actions and by the variability of advice given by child users, with its impact on the emotional state of the victim.

It is worth stressing that the character-based approach taken in our work, and in that just cited, is the converse of plot-centred approaches in which planning can also be used. In the plot-centred approach, a single planning system generates actions for characters to execute in an approach much closer to the classical planning paradigm, but at the cost of any character-based autonomy [28]. Reconciling this with character believability is non-trivial, but more seriously, while this approach supports the user as author, it leaves little or no scope for the user as participant to influence narrative unfolding on the fly, as in our work.

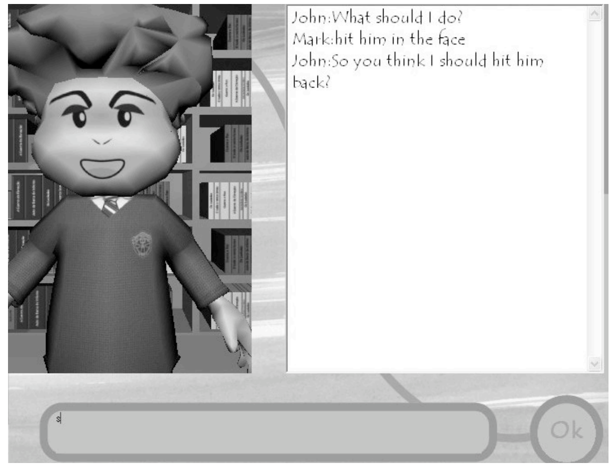

Figure 6: The interaction window: advice from a child

Conclusion and further work

In this paper we have discussed the integration of a planning system into a believable graphical character so as to support the generation of short dramatic episodes via character interaction. Where a conventional planner might focus on correctness and efficiency, the focus in an IVA is on responsiveness and believability. For this reason, it is not clear that the modern generation of graph-planning algorithms has anything very useful to contribute to the problem because their style of operation is monolithic, detached from execution and extremely difficult to make interactive.

Unlike many conventional planning domains, plans are typically short while goals are numerous and often relatively short-lived. Thus goal management is at least as important as planning alone, a point that has been made by other researchers who focus on applications of planning rather than on pure algorithm development [25].

Of crucial importance in the system described is the integration of an affective appraisal system with the planning mechanism. It is this that supports the flexible management of goals that is one of the requirements for the continuous planning strategy adopted and allows a differentiation in behaviour for different characters. To the extent that it is controlling the importance of goals, the affective system is also acting as a heuristic source of direction since it is also identifies goal conflicts that would not be registered in a more conventional planner.

The affective system adopted in this work operates at a purely cognitive level, and to that extent it may be argued it is somewhat implausible from a biological point of view: lower-level systems working through more primitive sensory channels have been investigated in other work and an open issue is how they might be integrated with the system that has been discussed. Alternative theories from psychology, such as the PSI theory of Dorner discussed in [5] are candidates for this type of integration.

The architecture has as yet only been applied to a physical bullying scenario in which pushing, hitting and kicking are as important as the language actions of mocking, insulting and threatening. Relational bullying – which typically involves girls rather than boys – is very much more language based, since it depends on social exclusion rather than physical intimidation for its effects. Work in the EU FP6 project eCIRCUS, due to start in early 2006, will develop the approach and examine whether the planning mechanism developed so far is adequate for this extension.

Acknowledgement

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

References

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

[2] Aylett R, Louchart S, 2003. *Towards a narrative theory of VR*. Virtual Reality Journal, Volume 7 January 2004

[3] Aylett, R.S (2004) Agents and affect: why embodied agents need affective systems Invited paper, 3rd Hellenic Conference on AI, Samos, May 2004 Springer Verlag LNAI 3025 pp496-504

[4]Aylett R, Louchart S, Pickering J (2004) A mechanism for acting and speaking for empathic agents. Workshop “Empathic Agents”, The Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS04), New York

[5] Bach, J. (2002) Enhancing Perception and Planning of Software Agents with Emotion and Acquired Hierarchical Categories. Proceedings, MASH02, German Conference on AI KI 2002, Karlsruhe, Germany pp3-11

[6] Bates, J. (1994), The Role of Emotion in Believable Agents Communications of the ACM v37, n7 pp122—125 1994

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

[8] Cassell, J., Bickmore, T., Billinghurst, M., Campbell, L., Chang, K., Vilhjálmsson, H. and Yan, H. (1999). "Embodiment in Conversational Interfaces: Rea." Proceedings of the CHI'99 Conference, pp. 520-527. Pittsburgh, PA

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

[10] Damasio, A. (1994)Descartes Error. Avon Books. 1994

[11] Dautenhahn,K (2000), Socially Intelligent Agents and The Primate Social Brain - Towards a Science of Social Minds, Proc. AAAI Fall Symposium "Socially Intelligent Agents - The Human in the Loop", AAAI Press, Technical Report FS-00-04, pp. 35-51.

[12] Dautenhahn, K (2003): Roles and Functions of Robots in Human Society - Implications from Research in Autism Therapy, Robotica, volume 21(4), pp. 443-452.

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

[14] Frijda, N. (1986). The Emotions. Studies in Emotion and Social Interaction. New York: Cambridge University Press

[15] Gratch,J. (1999) "Why you should buy an emotional planner," Proceedings of the Agents'99 Workshop on Emotion-based Agent Architectures (EBAA'99)

[16] Gratch J. (2000) Émile: Marshalling Passions in Training and Education. In 4th International Conference on Autonomous Agents, ACM Press, June 2000

[17] Lazarus, R.S& Folkman, S. (1984). Stress, appraisal and coping. New York: Springer

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

[19] Muller, J. P., Pishel, M. & Thiel, M. (1995), "Modelling Reactive Behaviour in Vertically Layered Agent Architectures", In Wooldridge, M. & Jennings, N. (eds.) (1995b), Intelligent Agents, Lecture Notes in Artificial Intelligence 890, Heidelberg: Springer Verlag, 261-276.

[20] Ortony,A, G. Clore, A. Collins. (1988) *The Cognitive Structure of Emotions*. Cambridge University Press

[21] Picard, R. 1997 *Affective Computing*. MIT Press.

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

[23] Sloman A.(1998) 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

[24] Http://www.VICTEC.org

[25] Pollack, M. & Horty, J.F (1999) Horty, There's More to Life Than Making Plans, AI Magazine20(4), pp. 71-83, 1999

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

[27] Aylett, R.S;. Louchart,S; Dias, J; Paiva, A and Vala, M. (2005) FearNot! - an experiment in emergent narrative. Proceedings, IVA 2006, LNAI 3661, Springer ,pp305-316 ISBN-10 3-540-28738-8

[28] Riedl, M & R. Michael Young (2004) An intent-driven planner for multi-agent story generation by. in the Proceedings of the 3rd International Conference on Autonomous Agents and Multi Agent Systems, July 2004.