Integrating Models of Personality and Emotions into Lifelike Characters

Elisabeth André, Martin Klesen, Patrick Gebhard, Steve Allen, and Thomas Rist

DFKI GmbH, Stuhlsatzenhausweg 3, D-66123 Saarbrücken, Germany,
{andre, klesen, gebhard, allen, rist}@dfki.de

Abstract. A growing number of research projects in academia and industry have recently started to develop lifelike agents as a new metaphor for highly personalised human-machine communication. A strong argument in favour of using such characters in the interface is the fact that they make human-computer interaction more enjoyable and allow for communication styles common in human-human dialogue. Our earlier work in this area has concentrated on the development of animated presenters that show, explain, and verbally comment textual and graphical output on a window-based interface. Even though first empirical studies have been very encouraging and revealed a strong affective impact of our Personas [23], they also suggest that simply embodying an interface agent is insufficient. To come across as believable, an agent needs to incorporate a deeper model of personality and emotions, and in particular directly connect these two concepts.

Introduction

The German Research Centre for Artificial Intelligence (DFKI) recently started three new projects to advance our understanding of the fundamental technology required to drive the social behaviour of agents. This initiative has been timed to catch the current wave of research and commercial interest in the field of lifelike characters (see [1]) and affective user interfaces (see [24]).

The Presence project uses an internal model of the agent's (and possibly the user's) affective state to guide the conversational dialogue between agent and user. The second project features an Inhabited Market Place in which personality traits are used to modify the character roles of virtual actors in interactive presentations. The i3-ese project Puppet promotes the idea of a virtual puppet theatre as an interactive learning environment to support the development of a child's emotional intelligence skills.

Although all three projects rely on a more or less similar approach towards modelling emotions and personality traits, there are variations with regard to the underlying user-agent(s) relationship(s), the factors that influence an agent's emotional state, and the way how emotions and personality traits are made observable. The following sections provide short overviews of the three projects and discuss their affective nature in more detail.

Basic Concepts

One of the first challenges we must face when attempting to use *affect* within our architectures, is to recognise the fact that the term does not refer to a well-defined class of phenomena clearly distinguishable from other mental and behavioural events. Affect is used within the literature to describe the class of motivational control states which result from valenced reactions to objects and events - these include emotions, mood, and arousal. Therefore the only generalisation we can really make about affect is that it must contain at least the two attributes of activation and valence. The different classes of *affective states* can further be differentiated by *duration*, *focus*, *intensity*, and *expression / effect* - emotions tend to be closely associated with a specific event or object and have a short duration, whereas mood is more diffuse and of longer duration. For the purposes of this paper, we will define personality as "the complex of characteristics that distinguishes an individual or a nation or group; especially the totality of an individual's behavioural and emotional characteristics", and emotion as "affect that interrupts and redirects attention (usually with accompanying arousal)" [21].

Although there is no consensus in the nature or meaning of *affect*, existing theories and models of personality and emotion can still play an useful role in enhancing user-agent interaction - even though they do not capture the *affective* phenomena in its entirety. As a starting point for our work, we have taken the Five Factor Model (FFM) [13] of personality, and the Cognitive Structure of Emotions model (OCC - Ortony, Clore and Collins) [15]. These models are readily amenable to the intentional stance, and so ideally suited to the task of creating concrete representations / models of personality and emotions with which to enhance the illusion of believability in computer characters.

Emotions: The OCC model of emotions provides a classification scheme for common emotion labels based on a valence reaction to events and objects in the light of agent Goals, Standards, and Attitudes. The OCC model is a model of causation, and will be used within both Presence and Puppet to determine the affective state of the character in response to events in the environment (see also [6] and [18]).

Personality: The FFM is a purely descriptive model, with the five dimensions (*Extraversion*, *Agreeableness*, *Conscientiousness*, *Neuroticism*, and *Openness*) being derived from a factor analysis of a large number of self- and peer reports on personality-relevant adjectives. The descriptive nature of the FFM gives us an explicit model of the character's personality, and in turn, allows us to concentrate on using the affective interface to directly express those traits (which offers the interesting possibility of attempting to recreate the character's personality traits from an analysis of the emergent social interaction). Furthermore, as we are focusing on social interactions, we can concentrate on the traits of *extraversion* (Sociable vs. misanthropic; Outgoing vs. introverted; Confidence vs. timidness) and *agreeableness* (Friendliness vs. indifference to others; A docile vs. hostile nature; Compliance vs. hostile non-compliance) - although we will also use *neuroticism* (Adjustment vs. anxiety; Level of emotional stability; Dependence vs. independence) to control the influence of emotions within our characters.

In addition to generating *affective states*, we must also express them in a manner easily interpretable to the user (which in the case of the Puppet project will be young children). Personality and emotions can be conveyed in various ways. According to

empirical studies, extravert characters use more direct and powerful phrases than introvert characters [8], speak louder and faster [20] and use more expansive gestures [9]. Furthermore, the rendering of dialogue acts depends on an agent's emotional state. Effective means of conveying a character's emotions include acoustic realisation, body gestures and facial expressions (see [5]). While these studies seem directly applicable to anthropomorphic agents like the Presence Persona, it is not clear to what extent they apply to animals with anthropomorphic features such as the characters in the Virtual Puppet theatre.

In all three projects, personality and emotions are used as filters to constrain the decision process when selecting and instantiating the agent's behaviour. For instance, we might define specific behaviours for extravert characters in a certain emotional state. However, there are other (affective) states we would like to convey that are not simply the result of an affective appraisal (as in the OCC model), or easily derived from personality traits - i.e. fatigue, boredom, and hunger. To model these states, we will mimic our character's active body state with motivational drive mechanisms to provide the affective input signals.

The Role of Affect in Presence

The Presence project will use lifelike characters as virtual receptionists / infotainers / accompanying guides for visitors to the German Research Centre for Artificial Intelligence (DFKI). Here we will explore the hypothesis that using an explicit affective model (of both agent and user) to guide the presentation strategies used in the human-agent conversational dialogue will (a) create a more natural and intuitive user interface (by tailoring the conversation to an individual person); (b) provide the user with an engaging and enjoyable experience; and (c) enhance the believability of virtual characters.

The Presence project addresses a number of specific problem areas: (a) flexible integration of multiple input (speech, mouse, keyboard and touch-screen) and output (text, pictures, videos and speech) devices. The architecture must be intelligent enough to adapt to the different affective modes of communication in the different application domains - i.e. no speech inflection in the remote domain; (b) the development of a high-level descriptive language for character definition, based on personality traits to allow easy customisation of the agent; (c) the combination of computational models of personality and emotion with planning techniques to guide the interaction of a lifelike character presenting material to visitors both locally and / or remotely over the world wide web; and (d) explore the possibility of tailoring the agent-user interaction to an individual user by inferring the user's affective state (see also [3]).

Application Domain

One of the major design considerations of the project is the requirement for a flexible architecture capable of coping with multiple application domains. Domains are defined by the scope of information they cover and the general interaction

behaviour of the lifelike character. Domains are furthermore decomposed into hierarchically structured dialog topics (i.e. Welcome/Small-talk, Helpdesk, Info/Tour) which guide the conversational thread. Our domains are:

- *Receptionist (Entrance Hall)*: Here the Presence system will run on an info-terminal within the DFKI entrance hall, and will welcome visitors, business partners, and student to the institute. The *virtual receptionist* will answer questions on a wide range of dialogue topics covering news, research projects, and people within the DFKI. The receptionist will also be capable of initiating conversations and informally introducing the various dialogue topics. The actual conversational thread will be guided by both the user responses, and the modelled affective state of the agent.
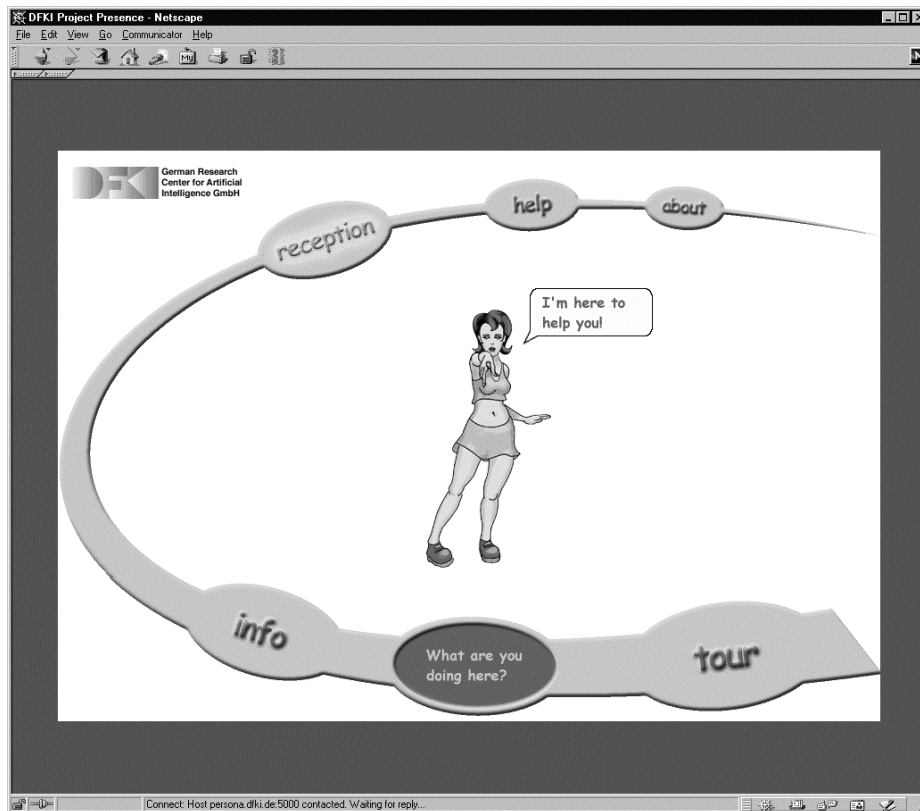

Fig. 1. Screenshot of the Presence Prototype.

- *Infotainment (Remote)*: The remote infotainment domain will utilise the same underlying conversational modes as the local receptionist domain. However, the constraints imposed by the restricted bandwidth of the internet will place an additional requirement on the system to make best use of the available input and output resources. Our system must therefore be intelligent enough to cope with a varying communication channel with the user, i.e. one in which the affective channel of speech inflection may no longer be available.
- *Guide (Portable)*: Within the portable application domain, the system's primary role will be to guide the user through the building, i.e. the user can ask the system how to reach a lab or office. However, instead of remaining a passive guide, the Persona system will take advantage of the infra-red channel of a palm computer to provide a low bandwidth link to the server - thus allowing the system to update the user with DFKI internal news, or to signal the beginning of lectures, meetings or talks. The portable domain will provide a real challenge to convey affective information in such an impoverished environment.

Emotions and Personality

The Persona system model extends the PPP animated presentation agent architecture developed at DFKI [2] with a number of new features - most notably enhanced input and output modalities for *affective* communication, and an *Affective Reasoning Engine* for *affective state* recognition and generation. In line with recent research in affective computing [4], [16] and [22], we use two *affective* information processing channels (see Fig. 2). *Primary* emotions (i.e. being startled, frozen with terror, or sexually stimulated) are generated using simple reactive heuristics, whereas *Secondary* emotions are generated by the deliberative *Affective Reasoning Engine* according to the OCC model – Sloman introduces the additional class of *Tertiary* emotions as *secondary* emotions which reduce self control, but these will not be implemented in our initial prototype.
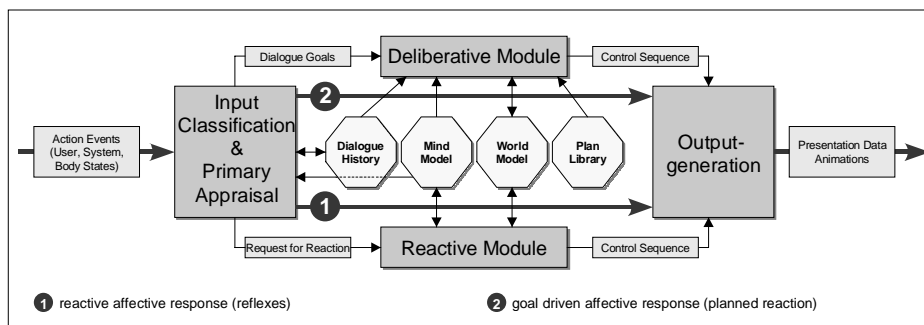

Fig. 2. Conceptual view of the Presence system which deals with affective states.

Emotions and personality traits weave an intricate web within the Persona system. Emotions are primarily used to determine the Persona's short-term affective state - expressed through the system's affective output channels as gestures and speech inflection. However, emotions are also able to directly influence the choice of phrase

within the conversational dialogue, and even dynamically adjust (within a pre-set range) the Persona's personality trait values. Likewise the Persona's personality traits: (a) help to bias the motivational profile of the agent - and thus determine the importance of the agent's high-level goals (to which events are compared during the emotion generation process); and (b) steer the conversational dialogue goals - extravert characters are more likely to initiate small-talk.

The Persona's *affective reasoning* process is roughly modelled on the "Emotion Process" described in [7]. The *Input Classification* and *Primary Appraisal* component classify incoming *action events* and appraise them with respect to agent concerns stored within the *mind model* and the *dialogue history*. After classification, filtered events are passed to the behaviour module as response requests through the two parallel affective information processing channels. The reactive module handles the Persona's immediate reactions to user or system events, whereas the deliberative module produces more controlled reactions (in particular, it is responsible for determining the contents and the structure of the dialogue contributions).

The rendering of the Persona's actions is performed by the *Output Generation* module which generates and co-ordinates speech, facial expressions and body gestures. So that the output generator can account for the Persona's affective state, the deliberative and reactive modules annotate the actions to be executed with appropriate mark-ups (in addition to personality traits and emotion label terms - i.e. happy, sad, fear - we will also use the general emotional dimensions of *Valence* and *Arousal*).

To model the Persona's personality, we use the social dimensions of *extraversion*, *agreeableness*, and *neuroticism*. We will initially model the Persona as an extravert, agreeable and emotionally-balanced character, as [11] point out, people tend to prefer others based on the match and mismatch to their own personality (even though the exact relationship is still unclear and empirical studies have led to conflicting results). Among other things, this means that the Persona will tend to take the initiative in a dialogue, will be co-operative and will remain patient if the user asks the same question over and over again (although the later case could indicate that the Persona is failing in her goal to be helpful).

We currently focus on goal-based emotions - whereby events are evaluated with respect to their desirability for the user's and/or the Persona's goals. We will also attempt to infer the user's affective state and use it to create a more sympathetic character. This can be done indirectly by monitoring system deficiencies, such as errors of the speech recognition component or the inaccessibility of information servers. Alternatively we can try to derive it directly from the syntactic and semantic form (use of affective phrases) and the acoustic realisation (talking speed, volume etc.) of the user's utterances.

The Role of Affect in The Inhabited Market Place

The objective of the Inhabited Market Place is to investigate sketches, given by a team of lifelike characters, as a new form of sales presentation. The basic idea is to communicate information by means of simulated dialogues that are observed by an audience. The purpose of this project is not to implement a more or less complete model of personality for characters, such as a seller and a customer. Rather, the

demonstration system has been designed as a test-bed for experimenting with various personalities and roles.

Application Domain

As suggested by the name, the inhabited market place is a virtual place in which seller agents provide product information to potential buyer agents. For the graphical realisation of the emerging sales dialogues, we use the Microsoft Agent package [14] that includes a programmable interface to four predefined characters: Genie, Robby, Peedy and Merlin. To enable experiments with different character settings, the user has the possibility of choosing three out of the four characters and assigning roles to them. For instance, he or she may have Merlin appear in the role of a seller or buyer. Furthermore, he or she may assign to each character certain preferences and interests (see Fig. 3).

Fig. 3. Dialog for character settings.

The system has two operating modes. In the first mode, the system (or a human author) chooses the appropriate character settings for an audience. The second mode allows the audience to test various character settings itself. Figure 4 shows a dialogue between Merlin as a car seller and Genie and Robby as buyers. Genie has uttered some concerns about the high running costs which Merlin tries to play down. From the point of view of the system, the presentation goal is to provide the observer – who is assumed to be the real customer - with facts about a certain car. However, the presentation is not just a mere enumeration of the plain facts about the car. Rather, the facts are presented along with an evaluation under consideration of the observer's interest profile.

Emotions and Personality

As in the other two projects, we follow a communication-theoretic approach and view the generation of simulated dialogues as a plan-based activity. We are investigating two approaches: a centralised approach in which the system acts as a screen writer who produces a script for the actors of a play; and a distributed approach in which the single agents have their own goals which they try to achieve.
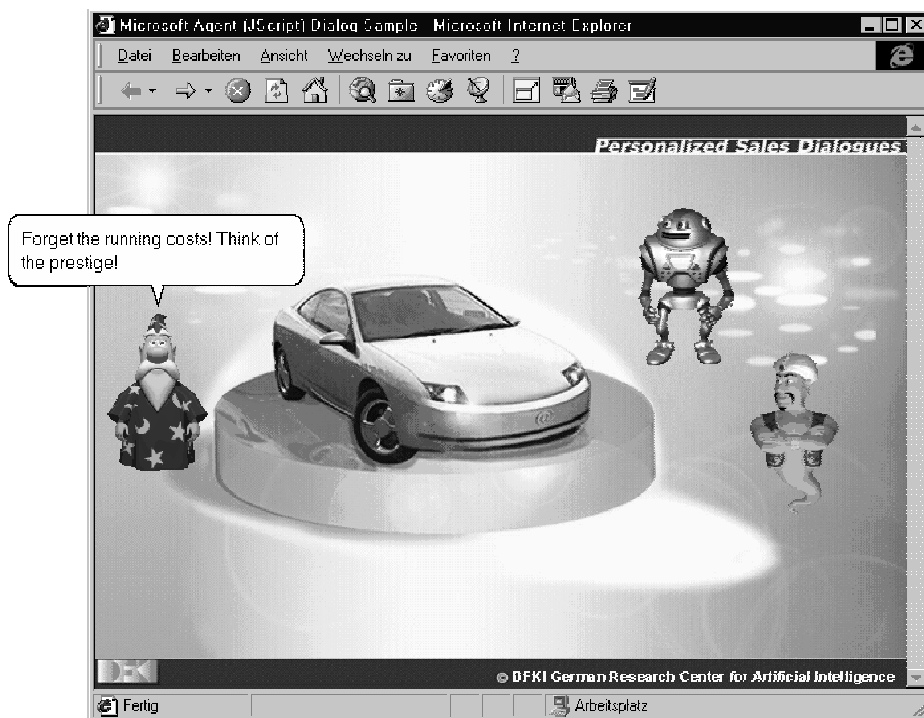

Fig. 4. Car sales dialogue example.

We follow the FFM, but model only the dimensions of extraversion and agreeableness which seem to be the most relevant dimension for social interaction (see [11]). In contrast to Presence, the Inhabited Market Place does not start from a fixed set of personality features, but allows the user to design the agents' personality him- or herself. We decided to model the same emotional dimensions as in the Presence project, namely *Valence* and *Arousal*.

In the Presence project, the agent addresses the user directly as if it were a face-to-face conversation between human beings. There is no direct communication between the user and the agents in the Inhabited Market Place. Consequently, an agent's emotional state is not influenced by the user, but by the behaviour of other artificial agents, and in particular by their dialogue contributions. To account for this, the speech acts of the dialogue participants are evaluated by the characters in terms of their role, personality traits and individual goals. The goals in particular determine the

desirability of events, for example, a buyer will be displeased if he is told that a relevant attribute of a car (e.g. power windows) is missing for a dimension that is important to him (e.g. comfort). In contrast to Presence, we do not explicitly model the user's emotional state. While the agent in the Presence project aims at positively influencing the user's emotional state, the agents in the Inhabited Market Place do not necessarily care for each other, i.e. an agent may intentionally make a dialogue contribution which has a negative impact on the affective state of another agent. Furthermore, the emotional responses of different agents to the same event may differ.

Personality in the Inhabited Market Place is essentially conveyed by the choice of dialogue acts and the semantic and syntactic structure of an utterance. Emotions in this scenario are expressed by facial expressions and the syntactic structure of an utterance. Since the Microsoft Agent Programming tool does not allow for intonational mark-ups, we do not convey emotions by acoustic realisation in this scenario.

First informal system tests with the Inhabited Market Place revealed some interesting aspects as to the importance of the matching of personality traits and surface characteristics. Subjects tended to believe that Merlin was more interested in comfort and safety while they expected that Robby was more interested in the technical details of a car. This also confirms our earlier empirical studies that showed that reusing the look and voice of characters for different roles is only possible to a certain extent.

The Role of Affect in Puppet

The objective of the Puppet project is to develop and investigate the value of a new virtual reality environment, the Virtual Puppet Theatre, (VPT), based on a theoretical framework of "learning through externalisation" [19]. Deploying user-controlled avatars and synthetic characters in the child's own play production, the children have to distinguish and master multiple roles in their interaction with the system, e.g. that of a director, an actor and an audience with the main activities producing, enacting and reflecting respectively. Within this process the children should gain a basic understanding on how different emotional states can change or modify a character's behaviour and how physical and verbal actions in social interaction can induce emotions in others. These emotional intelligence skills are important for us with respect to the early learning goals: "social role decentring" and theory of mind. Our approach is similar to [10] which allows children to direct a puppet's mood, actions and utterances in interactive story-making and to [12] where children may induce some changes in their characters emotional state besides selecting a characters actions.

Application Domain

For our first prototype (VPT1) developed for children at the age of 5-6, we decided to model a farmyard as a co-habited virtual world, in which the child's avatar (e.g. the farmer) and a set of synthetic characters (pigs, cows, etc.) can interact with each other. Figure 5 shows a screenshot of an early version using different drawings for each character, which will be transformed to full 3D models at a later stage. Our characters are designed to exhibit both physical and verbal behaviour. We do not try to model "real" animals but make them more cartoon-like instead.

Fig. 5. 2D Prototype of the farmyard scenario.

For the communication between the avatar and a character we will use a simple speech-act based dialogue model and a set of pre-recorded utterances. The agents are equipped with virtual sensors and effectors which connect them to the 3D virtual environment and controlled by an agent architecture that integrates deliberative (goal-driven) and reactive (data-driven) planning. To foster the above mentioned emotional skills we provide two distinct sets of interfaces which can be used by the child to control a character's behaviour. A *body control interface* which gives full control over the movement of the selected character and a *mind control interface* which allows to change the character's emotional state thus biasing the behaviour in some direction without specifying the actual motion pattern. Similar to the system described by [10]

we separate the high-level behaviour planning and affective reasoning (the "mind") from the animation planning and control modules (the "body"). The first is done by the agent architecture as described in the next section and the latter lies within the responsibility of the 3D virtual environment.

Emotions and Personality

The agent architecture used for the high-level behaviour planning and affective reasoning consists of a knowledge base, a plan library, an interpreter and an intention structure. The knowledge base is a database that contains the world model (the "beliefs") of a character. The plan library is a collection of plans that an agent can use to achieve its goals and the intention structure is an internal model of the current goals (the "desires") and instantiated plans (the "intentions") of that agent. Within this architecture we can have multiple active goals and multiple plans for each goal. Conflict resolution, plan selection, instantiation and execution are handled by the interpreter (see Fig. 6).
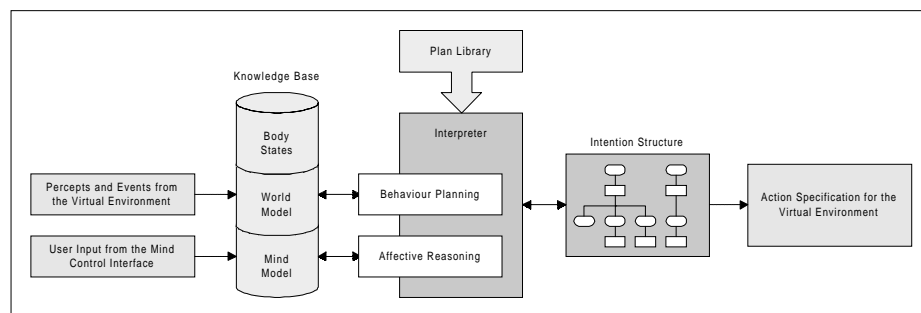

Fig. 6. Agent architecture for VPT1.

There are two types of input which can influence an agent's behaviour planning and affective reasoning: percepts and events from the virtual environment; and user input from the mind control interface. The first is used to update the agent's world model (which also influences the affective reasoning process) and the second to directly change its affective state encoded in the mind model. To increase the life-likeness of a character we also introduce body states (fatigue, boredom, hunger) which are represented within the knowledge base and regularly updated by the interpreter. The body states act as motivational drives that impel the agent into action by activating the appropriate behaviour (sleeping, playing, eating and drinking), which is then carried out in a character-specific way (if a dog is hungry it will go to the farmhouse whereas a cow will start grazing). The output of the planning processes is an action specification for the virtual environment. It contains appropriate mark-ups (e.g. for the facial expression) taking into account the current emotional state.

As a starting point for the mind model capturing the personality and affective states of our characters we use the OCC model. In Puppet it is particularly important that we can express these states in a manner easily interpretable by young children. We therefore decided to model the emotion types *Anger*, *Fear*, *Happiness* and *Sadness*

based on evidence suggesting their universality [5] and the fact that there are distinctive facial expressions which can be interpreted properly by children of age 4-8 [17]. Emotions are primarily conveyed by facial expressions and the selection of appropriate sounds (a cat will purr if it is happy and hiss if it is angry). They are either computed by emotion generating rules according to the OCC model or directly manipulated by the child through the mind control interface (see Figure 6).

As in the other two projects described in this document we adopt the FFM and reduce it to the dimensions extroversion and agreeableness because they determine to a large extent how an agent will behave in social interactions ([11]). In addition we specify for each character a set of preferences (e.g. the dog likes bones) and long term goals. Most characteristics are tailored for each character to give them unique pseudo personalities. This means that we can only partially rely on the same high-level behaviour to convey personality features (e.g. greet another character and start playing if you are extrovert and agreeable) and that we have to devise character-specific ones otherwise.

Puppet offers a variety of different user-agent(s) relationship(s). In "enacting mode" the child uses an avatar to interact with other characters in the scene. This is similar but not identical to Presence where the user interacts with the Persona through a set of input devices. The second mode, the child playing the role of an audience by observing the interaction of two or more autonomous agents has its equivalent in the Inhabited Market Place where the user observes a dialogue performed by a team of characters. However there is a third distinct user-agent relationship in Puppet, namely that of the child being a director, i.e. controlling the behaviour of *all* characters in the scene. This is similar to make-believe play with physical puppets during childhood in which the child takes on a series of roles. The difference is that the animals in our scenario are semi-autonomous, i.e. they take directions (e.g. the child can force the puppets to do or say something or change its affective states) that bias but not completely specify their behaviour. How an agent will (re)act in a specific situation also depends on its internal body states and personality features. This mode could provide valuable insights because we can observe when and how the children change the emotional state of a character, something that is not so easy to infer in conventional make-believe play.

The psychologists from Sussex will evaluate a group of 5-6 year old children before and after they played with the VPT. We hope that their findings will validate our assumption that the children's emotional intelligence skills were improved by constructing simple models of the virtual puppets minds. It will be also interesting to see how the ability to take the subjective view of different characters (as an actor) and to direct their behaviour (in the role of a director) will increase their understanding of the dynamics in social interactions, especially how emotions influence these interactions.

Conclusions

The three projects described in this paper use personality and emotions to emphasise different aspects of the affective agent-user interface: (a) *Presence* uses affect to enhance the believability of a virtual character and produce a more natural conversational manner; (b) the *Inhabited Market Place* uses affect to tailor the roles of actors in a virtual market place; and (c) *Puppet* uses affect to teach children how the different emotional states can change or modify a character's behaviour and how physical and verbal actions in social interactions can induce emotions in others.

As this broad range of application areas demonstrate, affect has an important role to play in user-agent interactions. However, affective user interfaces are still in their infancy, and much work is still needed to bring all the pieces of the jigsaw together. To use *affect* effectively, it must be an all-encompassing component of the system - from graphic design to system architecture to application contents.

Acknowledgements

References

1. André E. (1999). Applied Artificial Intelligence Journal, Special Double Issue on Animated Interface Agents, Vol. 13, No. 4-5.
2. André, E., Rist, T. and Müller, J. (1999). Employing AI Methods to Control the Behavior of Animated Interface Agents. Applied Artificial Intelligence 13:415-448.
3. Ball, G. and Breese, J. (1998). Emotion and Personality in a Conversational Character. Workshop on Embodied Conversational Characters. Oct. 12-15, Tahoe City, CA, pp. 83-84 and 119-121.
4. Damasio, A. R. (1996). Descartes' Error. London: Papermac. (first published 1994, New York: G. P. Putman's Sons.)
5. Ekman, P. (1992), An Argument for Basic Emotions, In: Stein, N.L., and Oatley, K. eds. Basic Emotions, p. 169-200, Hove, UK: Lawrence Erlbaum.

6. Elliott, C. 92 (1992). The Affective Reasoner: A process model of emotions in a multi-agent system. PhD Thesis, Northwestern University, Institute for the Learning Sciences Tech. Report #32.

7. Frijda, N. H. (1986). The Emotions. Cambridge: Cambridge University Press.

8. Furnham, A. (1990). Language and Personality. In: H. Giles and W.P. Robinson (eds.) Handbook of Language and Social Psychology. (pp. 73-95). Chichester, England UK: John Wiley & Sons.

9. Gallaher, P.E. (1992). Individual Differences in Nonverbal Behavior: Dimensions of Style. Journal of Personality and Social Psychology, 63(1): 133-145.

10. Barbara Hayes-Roth, and Robert van Gent (1997), Story-Making with Improvisational Puppets, Agent '97, Marina del Rey CA, USA.

11. Isbister, K. and Nass, C. (1998). Personality in Conversational Characters: Building Better Digital Interaction Partners Using Knowledge About Human Personality Preferences and Perceptions. Workshop on Embodied Conversational Characters. Oct. 12-15, Tahoe City, CA, pp. 103-111.

12. Isabel Machado, Ana Paiva (1999), Heroes, Villains, Magicians,...: Believable Characters in a Story Creation Environment, AI-ED '99 Workshop on Animated and Personified Pedagogical Agents, Le Mans, France.

13. McCrae, R. R., and John, O. P. (1992). An introduction to the five-factor model and its applications. Special Issue: The five-factor model: Issues and applications. Journal of Personality 60: 175-215, 1992.

14. Microsoft Agent: Software Development Kit, Microsoft Press, Redmond Washington, 1999. URL: http://microsoft.public.msagent

15. Ortony, A., Clore, G. L., and Collins, A. (1988). The Cognitive Structure of Emotions. Cambridge: Cambridge University Press.

16. Picard, R. W. (1997). Affective Computing. Cambridge, Mass: The MIT Press.

17. Reichenbach, I., and Masters, J. (1983), Children's use of expressive and contextual cues in judgements of emotions. Child Development, 54, p. 102-141.

18. Reilly, W. S. (1996). Believable Social and Emotional Agents. Ph.D. Thesis. Technical Report CMU-CS-96-138, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. May 1996.

19. Scaife, M., Rogers, Y. (1996), External Cognition: How Do Graphical Representations Work? International Journal of Human-Computer Studies, 45, p. 185-213.

20. Scherer, K.R (1979). Personality Markers in Speech. In: K.R. Scherer & H. Giles (eds.). Social Markers in Speech, Cambridge University Press: Cambridge, pp. 147-209.

21. Simon, H. (1967). Motivational and emotional controls of cognition. Reprinted in Models of Thoughts, Yale University Press, (1979), pages 29-38.

22. Sloman, A. (1999). Architectural Requirements for Human-like Agents Both Natural and Artificial. (What sorts of machines can love?). To appear in K. Dautenhahn (Ed.) Human Cognition And Social Agent Technology, John Benjamins Publishing.

23. van Mulken, S., André, E. and Müller, J (1998). The Persona Effect: How Substantial is it? In: Proc. of HCI'98, Sheffield, pp. 53-66.

24. Velásquez Juan D. (1999). Emotional Architectures, Agents 99 Workshop, Seattle, WA.