

Creating Adaptive Affective Autonomous NPCs

Mei Yii Lim¹, João Dias², Ruth Aylett¹, and Ana Paiva²

¹ School of Mathematical and Computer Sciences,
Heriot Watt University,
Edinburgh, EH14 4AS, Scotland
`M.Lim@hw.ac.uk`, `ruth@macs.hw.ac.uk`

² INESC-ID, Instituto Superior Técnico,
Lisboa, Portugal
`{joao.dias, ana.paiva}@gaips.inesc-id.pt`

Abstract. The paper reports work to create believable autonomous Non Player Characters in Video games in general and educational role play games in particular. It aims to increase their ability to respond appropriately to the player’s actions both cognitively and emotionally. We discuss the modeling of adaptive affective autonomous characters based on a biologically-inspired theory of human action regulation taking into account perception, motivation, emotions, memory, learning and planning. These agents populate an educational Role Playing Game, ORIENT (Overcoming Refugee Integration with Empathic Novel Technology) dealing with the cultural-awareness problem for children aged 13 to 14.

1 Introduction

About a decade ago, the importance of AI in the field of computer games rose sharply when it became a checklist item in game design and development [1]. This was marked by the appearance of games using artificial life (A-Life) techniques, for example the Sims³. Characters in the Sims have various motivations and needs with ‘smart terrain’ in the game environment guiding their actions in various ways to satisfy those needs. Producing more responsive and competent Non Player Characters (NPCs), for example in Computer Role Playing Games (CRPGs), would expand their role beyond merely that of tactical enemies. It would lead to the emergence of games where challenges stem from social interactions with these characters rather than only from competition.

NPCs in computer games vary in importance and may play roles as bystanders, allies of or competitors with the player in the game’s fictional world. Unfortunately, their behaviour is so far usually scripted and automatic, defined by finite state machines (FSMs) in which specific actions are triggered by certain actions of or dialogue with the player. The designer makes use of tricks to force players through essentially linear stories. While FSMs are a good way of specifying pre-determined sequential behaviour, they are insensitive to contextual

³ <http://thesims.ea.com>

change, so that in general an NPC will perform the same action in response to the player no matter what else has happened since the previous occurrence. This tends to produce gameplay that is repetitive and thus unnatural.

A few games, such as Blade Runner [2] incorporated a crude autonomy, with simple goals and dynamic scripting into their NPCs and story line, but these are extremely narrow goals supported by very limited sets of behaviours. A more advanced and recent example would be Fallout3⁴. Traditional methods - finite state machines and rule based systems - are still adopted to create these somewhat autonomous NPCs. However establishing the finite states and rules can be very tedious. The designer has to figure out all the possible branches of the game and the programmer must then implement the correct behaviour for the NPCs and the game world. The task is further complicated by the increasing complexity and scale of current CRPGs and a trend toward a more open-ended gameplay which could easily lead to an explosion in the number of states and rules required. As the size of states and rules sets increase, the debugging and maintenance tasks also become harder. Hence, it would be desirable for some of this burden to be offloaded from the designer's shoulder [3].

A way to achieve this is by adding 'real' autonomy to NPCs through planning thus requiring only specification of high-level goals while the characters' reasoning engine or planner takes care of low-level details to achieve these goals. An example of one of the earliest game to utilize planning is F.E.A.R. [3]. Relationships between goals and actions are established through symbolic representation of the goal states and action preconditions and effects without the need for explicit rules to handle different situations in the game. As explicit transitions specification is not required, this method generates a more comprehensible representation compared to the traditional approaches, particularly when the number of game branches increases. The decoupling of goals and actions allows different NPCs to satisfy goals in different ways and allows layering of simple behaviours to create complex behaviour. Using this approach, NPCs are able to dynamically react to player actions and pro-actively make decisions in reasonable and realistic ways. This produces more intelligent behaviour while designers now only have to consider *what* the NPCs can do, not *when* and *how* NPCs decide to do it.

In order to establish a natural interaction between NPCs and a player, NPCs should behave adaptively and believably. Natural interaction is very important because it transforms the challenge of a game from a technical one to an interpersonal one, and thus may increase both the enjoyment and the engagement of players. The focus of the human perception of believability is on the social and emotional dimension of these NPCs. Bates [4] claimed that the greatest significant quality to believability in computer characters is appropriately timed and clearly expressed emotions. This is because emotions play a critical role in processes such as rational decision-making, perception, human interaction, human creativity and human intelligence [5]. Characters that are able to express their feelings and can react emotionally to events are more 'life-like' and are

⁴ <http://fallout.bethsoft.com/eng/home/home.php>

more likely to create a ‘suspension of disbelief’ [6] in players. *“The question is not whether intelligent machines can have any emotions, but whether machines can be intelligent without any emotions”* [7]. Furthermore, Damasio [8] provides neurological support for the idea that there is no ‘pure reason’ in the healthy human brain but emotions are vital for healthy rational human thinking and behaviour.

However, emotion alone is insufficient if the resulting NPCs do not learn from experience and do not adapt to its environmental circumstances or the player’s actions. In order to survive in a dynamic environment, NPCs need to cope with uncertainty, react to unanticipated events and recover dynamically in the case of poor decisions. An action regulation mechanism is needed that drives the NPCs’ behaviour and emotions so that it behaves in ways a human might expect. Plausible and consistent behaviour create an impression of personality in the NPC [9] and according to the famous Bugs Bunny animator, Chuck Jones, personality is what gives a character life [10].

A motivational system, the most basic system for maintaining internal equilibrium in humans and animals can provide the required mechanism. Although motivations and needs were employed in early A-Life games, they were not combined with more advanced techniques such as planning or emotional control systems. Some examples of autonomous agents with emotions can be found in the area of education [11, 12], social simulation [13, 14] and therapy [15]. However, up-to-date, affective autonomous agents that are capable of improvisational actions, appear to be able to ‘think’, and have desires, motivations, emotions, personality and goals of their own, are still rare in computer games. In this paper, we discuss our approach - combining motivation, emotion and cognition - to create adaptive affective autonomous NPCs that are designed to interact socially with players. The main aim is to create NPCs that are able to learn from experience, at the same time exhibit behavioural variation through adaptation to different environmental circumstances and its current needs. We start by a short discussion on educational role playing game. This is followed by a review of some related work. Next we describe in detail the ORIENT agent mind architecture by first mentioning the inspirations to the architecture and then exposing the implementation details. Section 5 presents an interaction simulation using the proposed architecture and Section 6 concludes the paper.

2 Educational Role Playing Games

Researchers point out that play is a primary socialization and learning mechanism common to all human cultures and many animal species. *“Lions do not learn to hunt through direct instruction but through modelling and play”* [16]. Serious games are effective because learning takes place within a meaningful context where what must be learned is directly related to the environment in which learning and demonstration take place.

According to Piaget and Lewin [17, 18], the social interaction that takes place in educational role-play (RP) acts as a stimulus for changing existing

beliefs and behaviour. Role-play supports the creation of knowledge and meaning from concrete - though imagined - experiences [19] and uses social interaction and emotional engagement as mechanisms for a learner-centred constructivist approach. Thus, cultural studies can be made exciting through an educational role play game. For instance, one in which the student is a space command member who must master the patterns of behaviour of an alien culture and pass as their friend within a digitally simulated world. The students will have interesting missions to keep them motivated and engaged. This approach shifts the students' cognitive effort from reading about educational content to hands-on experience of achieving compelling goals. Members of a team can cooperate with each other to solve the team's conflicts with other agents, whether a player from another team or an NPC. Such an opponent must be perceivable as endowed with a personality, appear believable and behave appropriately if the player is to be able to suspend disbelief in the way engagement with the storyworld requires. This in turn will lead to an enhanced learning experience [20].

In ORIENT⁵, our game world is designed with just such concept in mind. It is an interactive computer assisted role-playing game where three players act as visitors to a foreign planet, ORIENT, that is inhabited by an alien race, *Sprytes*, as shown in Figure 1. The main aim of the game is to promote cultural-awareness and integration of refugee/immigrant children in schools. It enables social and emotional learning in a secure social setting - virtual environment - and employs immersive devices as the interaction modalities [21]. The players' mission is to save the planet from a meteorite that is on its destruction course. In order to achieve this goal, the players have to cooperate with each other and the alien inhabitants, which means integrating themselves into the Spryte's culture [22] which is based on Hofstede's cultural dimensions [23]. They will have to learn to appreciate the cultural differences and exhibit acceptable social behaviour to gain the Sprytes' trust. The users will witness the Sprytes eating habits, life cycles, educational styles, family formation and value system [24] during the game. Since the game incorporates a social setting, each NPC must be able to establish social relationships with other NPCs and the players to ensure successful collaboration.

The ability to empathize, that is, to detect the internal states of others and to share their experience, is vital to the formation of long-term relationships [25] which in our case refers to social relationships that involve repeated interactions between two entities. According to Hogan [26] “...*empathy means the intellectual or imaginative apprehension of another's condition or state of mind without actually experiencing that person's feelings...*” (cognitive empathy), whereas Hoffman [27] posits that “...*empathy [is] a vicarious affective response to others...*” (affective empathy). Since enhancement of integration in a cultural group relies both on the understanding of the internal states of the persons involved and their affective engagement, both cognitive and affective empathy are relevant in ORIENT. Additionally, previous experience is crucial in maintaining long-term social relationships, which means a requirement for an autobiographic memory

⁵ <http://www.e-circus.org/>

[28] is inevitable. Autobiographic memory stores significant context-based information about previous encounters, allowing rich recollective experience, useful as a guide to future interaction. Through an ability to retrieve relevant previous experiences from its autobiographic memory, an NPC will be able to know how to react or adapt sensibly to a similar future situation. Thus, ORIENT provides a good case study for modelling NPCs with adaptive and improvisational capabilities, that possess autobiographical memory, individual personality and show empathy.



Fig. 1. Sprytes in ORIENT

3 Related work

Much work has been carried out on developing agents with autonomous capabilities (e.g. [11, 12, 13, 14, 15]). As discussed above, emotion plays an important role, acting as an evaluation mechanism on performance, filtering relevant data from noisy sources, and providing a global management over other cognitive capabilities and processes. These are significant capabilities when operating in complex real environments [29]. New emotional models are still being developed, therefore to explore all of them here would be tedious, if not impossible. Hence, only architectures that are relevant to this research are reviewed. Some of this work focuses on more physiological aspects while some focuses instead on cognitive aspects of human action regulation.

Examples of existing physiological architectures are those by Cañamero [30], Velásquez [31], Blumberg [32] and Pezzulo and Calvi [33]. Cañamero’s architecture relies on both motivations and emotions to perform behaviour selection for an autonomous creature. Velásquez developed a comprehensive architecture of emotion based on Izard’s four systems model [34], focusing on the neural mechanism underlying emotional processing. Blumberg developed an animated dog, Silas, that has a simple mechanism of action-selection and learning combining

the perspectives of ethology and classical animation. A more recent implementation of the model is AlphaWolf [13], capturing a subset of the social behaviour of wild wolves. Pezzulo and Calvi implemented an agent architecture focusing on modulatory influences of motivations on behaviour including a sensorimotor system - perceptual and motor schemas; and a motivational system - drives. All these architectures are useful for developing agents that have only existential needs and behave reactively, but are too low level for characters which require planning and storytelling capabilities as in ORIENT. Another problem of these architectures is that the resulting agents do not show emotional responses to novel situations because all behaviours are hard-coded.

An important feature neglected by the physiological approaches is that emotions involve evaluations. The specific nature of the individual’s emotion is a function of their appraisal of the situation as having some significance to themselves. The concept of appraisal was first introduced by Arnold [35]. She defined emotions as ‘felt action tendencies’ that characterise experience and differentiate it from mere feelings of pleasantness or unpleasantness. The OCC cognitive theory of emotions [36] is one of the most used emotion appraisal models in current emotion synthesis systems. The authors view emotions as valenced reactions that result from three types of subjective appraisals: the appraisal of the desirability of events with respect to the agent’s goals, the appraisal of the praiseworthiness of the actions of the agent or another agent with respect to a set of standards for behaviour, and the appraisal of the appealingness of objects with respect to the attitudes of the agent. Numerous implementations of the theory exist, aimed at producing agents with a broad set of capabilities, including goal-directed and reactive behaviour, emotional state and social knowledge. These begin with the Affective Reasoner architecture [37] and the Em component [38] of the Hap architecture [39], with later implementations including EMA [40], FAtiMA (FearNot! Affective Mind Architecture) [41], and many more.

In order to create purely autonomous agents, we argue that a hybrid architecture combining both physiological and cognitive aspects is required. The agent cognitive processes should result from lower-level physiological processing and the outcome of cognitive processes should influence the agent’s bodily states, producing complex behaviours that can be termed emotional. This is supported by Damasio’s proposal of the existence of a body-mind loop in emotional situations [8]. Opposing appraisal theories, Izard [34] provided evidence that using cognitive processes alone to explain emotion activation is incomplete. He took cognitive processes as one of several factors that influence emotion generation rather than as a necessary or sufficient factor. Affirming this view, psychologist Dietrich Dörner proposed the ‘Psi’ theory [42, 43], integrating cognition, emotion and motivation for human action regulation. The ‘Psi’ theory is based on the argument that humans are motivated emotional-cognitive beings. Some other examples of this type of architecture are those by Sloman [44], Jones [45] and Oliveira [29].

At the center of deliberative agent architectures lies the BDI (Beliefs, Desires, Intentions) model [46] that has its roots in the theory of practical reasoning,

involving two main processes: deliberation (deciding what states of affairs we want to achieve) and means-ends reasoning (deciding how to achieve these states of affairs). This architecture provides a separation between plan selection and execution allowing BDI agents to balance the time spent deliberating about plans and executing those plans. A BDI agent derives intentions from its beliefs and desires. It then selects a set of actions to achieve these intentions. Thus, desires are the main source of behaviour but this in itself offers no ground rules to help authors in designing agents with believable behaviour.

The problem of the BDI model is that it is only concerned with the connection between deliberation and means-ends reasoning to achieve practical reasoning, and it does not provide any theoretical ground to answer question such as: What should be an agent's desires (and why)? Why should a desired be preferred over another?. These questions correspond to the functions *options* and *filter* which are clearly unspecified in the BDI model. Additionally, BDI agents usually do not learn from experience or adapt to their interaction environment. These issues must be overcome because the reason why a character pursues a given goal is as important as achieving the goal if we want the viewers to understand that the actions of the character are driven by its internal state [6].

4 ORIENT Agent Mind

4.1 Inspiration

FAtiMA

The ORIENT agent mind (i.e. the program that controls the behaviour of its NPCs) is built upon FAtiMA [41] architecture applied in FearNot!v2.0. FAtiMA was developed with the goal of creating believable synthetic characters, where emotions and personality play a major role in the character's decision making processes and are visible through the character's actions. FAtiMA was an extension of a BDI architecture and has a reactive and a deliberative appraisal and action selection layer as shown in Figure 2.

The reactive appraisal process matches events⁶ with a set of predefined emotional reaction rules providing a fast mechanism to appraise and react to a certain situation. An emotional reaction rule associates an external event to OCC's appraisal variables: desirability of the event, desirability for others and praiseworthiness of the action. The exact value of the variables for each rule is authored beforehand. These variables are then used according to the OCC theory to generate a wide range of emotions from Joy and Distress to Pity and Anger. The deliberative appraisal layer generates emotions by looking at the state of current intentions, more concretely whether an intention was achieved or failed, or the likelihood of success or failure. Together, the deliberative and reactive appraisal handle all of OCC's 22 emotions.

⁶ Events represent external actions that happen in the virtual world. Abbuk Greet-Gesture User is an example of such an event.

After the appraisal phase, both reactive and deliberative components perform practical reasoning. The reactive layer uses simple and fast rules that directly associate emotions to single actions. These rules, named action tendencies, define a character's impulsive reactions to particular emotional states (e.g. crying when very distressed, kicking an object when angry), that he does without thinking. On the other hand, the deliberative layer uses BDI-like deliberation and means-ends reasoning processes, hence takes longer to react but allows a more sequentially complex and goal-oriented behaviour.

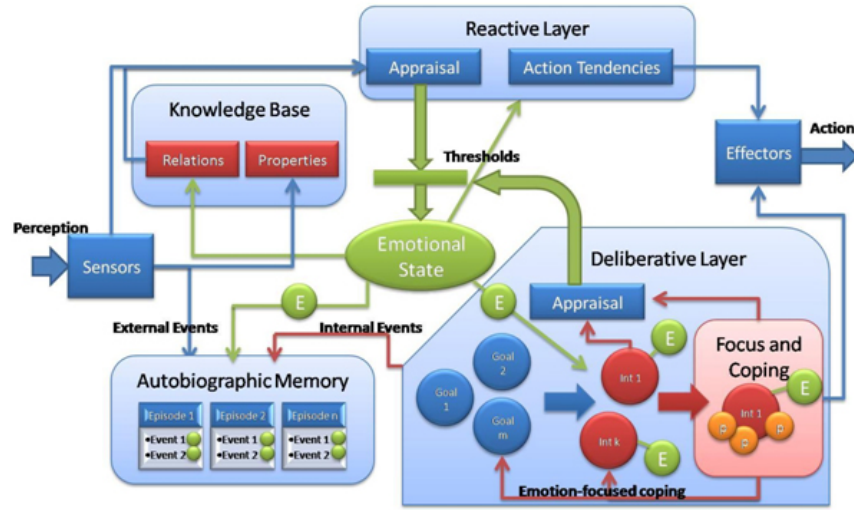


Fig. 2. FAtiMA architecture

The deliberative layer follows deliberation and practical reasoning processes as structured by the BDI model. All behaviour is driven by defined goals, which correspond to desires in BDI terminology. There is one additional difference from the generic BDI model in which intentions are only created after the filtering function. Once a goal becomes active, an intention to achieve such goal is created immediately (representing a commitment to achieve the goal) and the filtering process just decides which of the committed intentions will be selected first for execution.

FAtiMA uses two types of OCC goals: the active-pursuit goal and the interest goal. Active-pursuit goals represent goals that the character actively pursues to achieve a certain state (e.g. going to the cinema to watch a movie) while interest goals are goals that the character continuously maintains to avoid threatening situations (e.g. to stay healthy).

An active-pursuit goal is activated only if its start conditions are all satisfied and contains the following attributes.

Table 1. Active-Pursuit Goal attributes

Attributes	Description
ID	The goal identifier or name
Preconditions	A conjunction of conditions that determines when the goal becomes active
SuccessConditions	A conjunction of conditions used to determine if the goal is successful
FailureConditions	A disjunction of conditions that determines the goal failure
ImportanceOfSuccess	Specifies the goal’s importance of success
ImportanceOfFailure	Specifies the goal’s importance of failure

The success conditions represent the world state that the agent desires to attain. If all these conditions are verified, the goal is considered achieved. If any of the goal’s failure conditions become true while the goal is active, it automatically fails and the intention to achieve this goal is removed. OCC distinguishes between the importance of the success of a goal and the importance of its failure. Suppose that a character has a goal of parachuting; if it thinks that it will be successful in performing the goal, it will not feel great satisfaction (low ImportanceOfSuccess), but if the character thinks it will fail to parachute, it will feel tremendously afraid (high ImportanceOfFailure). Both these variables are important to achieve different intensities (Hope and Fear) for the same goal. The strength of the emotions are used to select between competing goals for planning and execution, the rationale being that the goals generating the strongest emotions are the most relevant ones.

An interest goal on the other hand has attributes as listed in Table 2.

Table 2. Interest Goal attributes

Attributes	Description
ID	The goal identifier or name
ProtectedConditions	A conjunction of conditions that must be preserved
ImportanceOfSuccess	Specifies the goal’s importance of success
ImportanceOfFailure	Specifies the goal’s importance of failure

FAtiMA has a a continuous planner [47] that is capable of partial order planning and includes both problem-focused and emotion-focused coping [48] in plan execution. Additionally, it appraises events to generate emotions and this information is stored in the autobiographic memory [49] for future reference. Problem-focused coping occurs when a character acts on the environment to tackle a situation. It involves planning a set of actions that achieve a desired result and executing those actions. On the other hand, emotion-focused coping works by changing the agent’s interpretation of circumstances, that is, lowering strong negative emotions for example, by lowering the importance of goals, a

coping strategy used often by humans when problem focused coping has a low chance of success.

The main reason for choosing FAtiMA over other agent architectures such as EMA or ALMA[], is that it incorporates the OCC theory [36] together with a continuous planner with problem-focused and emotion-focused coping. The OCC model includes emotions that concern behavioural standards and social relationships based on like/dislike, praiseworthiness and desirability for others, and thus it allows appraisal processes that take into consideration cultural and social aspects, very important for ORIENT agents [22, 50]. Additional technical reasons lead us to prefer FAtiMA. For instance, FAtiMA was already integrated with the graphical engine used to build ORIENT, OGRE3D ⁷, which allowed for a faster development of the application.

Another advantage of OCC theory is that it provides us with a simple means for modelling empathic emotions by taking into consideration appraisals of events regarding the consequences for others. It is - as far as we know - the only model that provides a formal description of non-parallel affective empathic outcomes (i.e. emotions that take a bad relationship between one agent and another into account, e.g. gloating and resentment). However, the number of empathic emotional outcomes described in OCC: happy-for, resentment, gloating and pity is limited.

On the negative side, FAtiMA does not take the physiological aspects of emotion into account. Furthermore, since it follows closely the BDI model and does not extend it with additional concepts that justify the appearance of goals/desires, FAtiMA faces the same problems of the BDI model.

Another problem with FAtiMA is the tedious authoring process of the character's goals, emotional reactions, actions and effects, and action tendencies so that the final behaviour of the characters is as intended. This is a difficult process because it is not easy to understand the interaction between distinct components such as goals and actions; and there is usually no theoretical ground for defining values such as the importance of success/failure of goals and desirability of events, which means that authors have to assign values that seem reasonable and adjust them by trying out the agent's behaviour and seeing if it corresponds to what they want. The character's personality is defined by a set of goals, a set of emotional rules, the character's action tendencies, emotional thresholds and decay rates for each of the OCC emotion types [41]. Because of this huge set of interacting factors, it is tedious to achieve a certain personality for a character.

Moreover, having these values scripted reduces the dynamism of some of the core aspects modeled, resulting in agents that are not adaptive and do not learn from experience. The character will always have the same goals and the same reactions regardless of the situation it is in. For example, eating is always desirable no matter whether the character has just eaten or not. If all characters behaved in the same way they would surely not be believable. Therefore, we looked for a theoretical framework that would help us achieve some of these

⁷ <http://www.ogre3d.org/>

behaviours without explicit authoring and a feedback mechanism that allows our agents to adapt to different situations.

PSI

We considered the PSI theory [42], a psychologically founded theory that incorporates all the basic components of human action regulation: perception, motivation, cognition, memory, learning and emotions, as shown in Figure 3. It allows for modelling autonomous agents that adapt their internal representations to a dynamic environment. Three successes of the PSI theory in replicating human behaviour in complex tasks can be found in [51, 42, 43].

A PSI agent does not require any executive structure that controls behaviour, rather, processes are self-regulatory and run in parallel, driven by needs. A deviation from the threshold on a need will give rise to an intention. Intentions are committed goals that afford the satisfaction of these needs. Thus, all action produced by a PSI agent is based on a limited number of basic needs modelled through homeostatic variables including

- existence-preserving needs (survival needs) - food, water and maintenance of physical integrity. These drives are relieved by the consumption of matching resources and increase by metabolic processes (in the case of food and water) or inflicted damage (in the case of integrity)
- species-preserving need - sexuality. This drive reflects the need of the organism to reproduce itself.
- need for affiliation - need for social experiences, e.g. to belong to a group or to be accepted by others. It increases with anti-legitimacy signals such as signs of social exclusion.
- need for certainty - being able to predict what will happen in a certain situation and also being able to predict consequences of (one's own) actions. The drive is achieved through exploration (to increase knowledge about the environment, hence improve prediction) and frustrated when outcomes mismatch expectations.
- need for competence - being able to master problems and tasks, including the ability to satisfy one's needs. Achievement of goals increases competence, while failure decreases it.

More than one need may be activated at a particular time leading to the activation of a few intentions with only one chosen for execution. As opposed to the BDI architecture, the PSI model provides information about its decision choices. An intention is selected for execution based on its anticipated probability of success, the degree to which it satisfies needs and its estimated urgency. If the character does not have knowledge on how to satisfy a goal, the success probability will be low, however, if its competence is high, it will perform explorative behaviour and may consider selecting the goal. PSI agent has three stages for intention execution. First, the agent tries to recall an automatic, highly ritualized reaction to handle the intention. If this is not successful or if no such

patterns that resemble emotional experiences in biological agents. For example, if an event leads to a drop in the character’s certainty and competence, then its *arousal* level increases causing a decrease in the *resolution level* and an increase in *selection threshold* value. In such situation, a quick reaction is required hence a time-consuming search is not selected. The character will concentrate on the task in order to stabilise the deviated need(s) and hence may choose to carry out the first action that it found feasible. The character may be diagnosed as experiencing anxiety due to its high needs, quick reaction and inaccurate perception and planning - prediction of emotion through recognised behaviour. Therefore, depending on the cognitive resources and the motivational state of the agent in a given situation, these parameters are adjusted, resulting in more or less careful or impulsive ways of acting, as well as more or less deliberate cognitive processing. In all processes, memory functions as a central basis for coordination.

Since FAtiMA already includes perception, cognition, memory and emotions, we added the PSI motivational and learning components into the existing architecture. The PSI motivational system serves as a quick adaptation mechanism of the agent to a specific situation and may lead to a change of belief about another agent as shown in [52], important for conflict resolution among ORIENT characters. PSI’s other advantage over FAtiMA is that it does not require much authoring except initialising the agents with some prior knowledge. Effects of actions are learned by trial and error and highly learned behaviours are reinforced by repeated use and satisfaction of needs. By trying different goals and actions under different circumstances, the agent will learn which goal and action is the most effective in satisfying its needs. PSI agents’ differences in behaviour will then correspond to different life-experiences that lead to different learned associations. For example, in FAtiMA, for each action, an emotional reaction rule has to be written for each agent to define the praiseworthiness of the action and the desirability (or undesirability) of the action to the agent itself and to other agents. Applying PSI on the other hand allows desirability (or undesirability) of events to be derived automatically from needs - the better an action or goal satisfies need(s), the more desirable it is. This involves slight modification to the way goals and actions are authored by including the potential effects on needs as a result of carrying out the corresponding goals or actions (explained in Section 4.2) but eliminates the emotional reaction rules sets completely. The same applies to interest goals and action tendencies rules set. Each agent in FAtiMA has a set of rules that specify its reactive actions for different situations, e.g. run when in danger. With PSI, this action is automatic because in such case, the need for integrity would be high leading the agent to choose the run action. Thus, PSI permits more flexibility both in authoring and the characters’ behaviour (adaptive) than FAtiMA. The reduction in authoring may not be obvious in the case of one agent with a few goals and actions but becomes more prominent when the number of agents, goals and actions increase.

Unfortunately, this very flexibility also leads to a lack of control over the characters’ behaviour. This is a problem because characters in ORIENT need to behave in certain ways so that the educational goals can be reached. This

problem becomes more visible when these characters must interact with users (e.g. perhaps ignoring the user because its need for affiliation is low). According to Squire [53], good educational games are games where narrative events situate the activity, constraining actions, provoking thought and sparking emotional responses. By making the NPCs react in certain ways, the player’s ability to access information or manipulate the world is limited. This forces the player to evaluate the relative value of information and to devise appropriate goals and strategies to resolve complex problems and help them to develop an experiential understanding of what might be otherwise an abstract principle. Therefore, there needs to be a balance between characters’ authored and learned behaviour. In the case of ORIENT, this would mean authoring possible goals and actions to avoid unpredictable behaviour but leave the characters to learn the effectiveness of goals and actions in satisfying its needs under different situations through trial and error. Coming back to the eating example (Section 4.1), a PSI agent will learn that eating when it is hungry is more desirable than eating when it is full, hence allowing it to select the best action according to its circumstances.

The advantages of combining FAtiMA and PSI

Combining FAtiMA and PSI, the problems of both psychological plausibility and control are addressed, neither of which can be solved by either architecture alone. Integrating the concept of drives into the BDI model helps us answer the two questions: What should be an agent’s desires? Why should a desired be preferred over another? In this new conceptual framework, desires correspond to situation states that the agent would ideally like to bring about in order to achieve one or more of the agent’s drives. All intentions and actions performed ultimately serve to achieve basic needs, and deliberation corresponds to selecting the desires that maximize the fulfilling of current needs.

Furthermore, cultural and social aspects of interaction can be modelled using FAtiMA as described in [22, 50] while PSI provides an adaptive mechanism for action regulation, fulfilling the requirements of ORIENT characters. Author are free to decide how much information they want to provide the characters to start with and leave the rest for the characters to learn. The degree of desirability (or undesirability) of an action or event is proportionate to the degree of positive (or negative) changes that an action or event brings to the agent’s drives. This desirability value can be used to automatically generate emotions according to the OCC model, removing the need to write predefined domain-specific emotional reaction rules. The other variables such as praiseworthiness and desirability for other in the emotional reaction rules are not addressed by this mechanism, but this problem is tackled through cultural parameterisation [22, 50]. Authoring of action tendencies and interest goals is also eliminated because now the characters always try to satisfy their needs which means avoiding threatening effects to their well-being.

Besides that, a simpler mechanism can be used to specify a character’s personality as discussed in the next section - specification of weights for 5 needs




Fig. 4. Situating the integration of FAtiMA and PSI in the Authored vs Learned dimensions

compared to specification of thresholds for all OCC emotions. Thus, the architecture achieves a balance between learning and authoring the characters' behaviour as shown in Figure 4. The resulting characters will seem biologically plausible, adaptive, have motivations, emotions, personality and goals of their own as well as being able to interact socially with players.

4.2 FAtiMA-PSI Architecture

In the ORIENT agent mind architecture shown in Figure 5, goals are guided by drives. Five basic drives from PSI are modelled in ORIENT including Energy, Integrity, Affiliation, Certainty and Competence. These drives can emerge over time or can be activated by events happening in the environment. Both Energy and Integrity are subsets of the existence-preserving need. Energy represents an overall need to preserve the existence of the agent (food and water). As the agent carries out actions, it consumes energy which means eventually, it will have to rest or perform actions to regain energy. Integrity represents well being, i.e. the agent avoids pain or physical damage while affiliation is useful for social relationships. On the other hand, certainty and competence influence cognitive processes. Note that we do not model species-preserving (sexuality) need since gender is absent from the Sprytes and this need is important only when interacting with a different gender.

Energy, Integrity and Affiliation drives

Table 3. FAtiMA-PSI Goal attributes

Attributes	Description
ID	The goal identifier or name
PreConditions	A conjunction of conditions that determines when the goal becomes active
SuccessConditions	A conjunction of conditions used to determine if the goal is successful
FailureConditions	A disjunction of conditions that determines the goal failure
EffectsOnDrives	Specifies the effects that the goal will have on the agent's or another agent's energy, integrity and affiliation drives if the goal succeeds

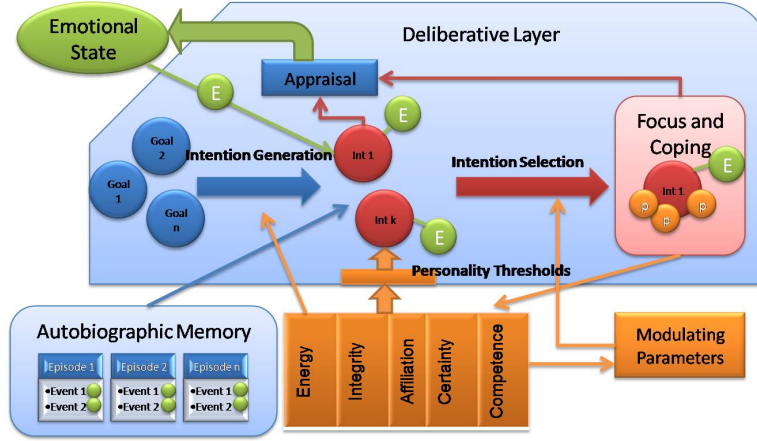


Fig. 5. FAtiMA-PSI architecture

In the FAtiMA-PSI architecture, a goal is defined by attributes listed in Table 3. Each goal contains information about its expected contribution to energy, integrity and affiliation, that is, how much the drives may be deviated from or satisfied if the goal is performed. It is not necessary to specify the contribution to certainty and competence since these are determined automatically as described later. Likewise, events or actions also include contributions to energy, integrity and affiliation drives. Based on this information, the importance of goals to each character at a particular time instance can be determined, allowing the character to give priority to goals that satisfy its deviated drives under different circumstances. This is an advantage over the previous FAtiMA architecture where a goal's importance is pre-authored which means that whenever a goal activation condition becomes true, the goal is always created with the same importance of success and failure, independently of the situation that originated the goal. This causes a problem in deciding which goal should be selected when there are several conflicting goals.

It is assumed that the scales for all drives are comparable, ranging from 0 to 10 where 0 means complete deprivation while 10 means complete satisfaction. An agent's aim is to maintain these drives at the highest level possible at all times in order to function properly. Each drive has a specific weight (discussed in Section 4.2) ranging from 0 to 1 that underlines its importance to the agent. Every time an action is executed, the agent's drives are updated according to the following equation:

$$Level_t(d, a) = Max(0, Min(10, Level_{t-1}(d) + (effect(a, d) * Weight(d)))) \quad (1)$$

The level of the drive d after action a is determined by the level of the drive at the previous instant ($Level_{t-1}(d)$) plus the effect that action a has on drive d ($effect(a, d)$). Effects of actions in drives are fixed and predefined for a given domain. The maximum and minimum functions are used to ensure that a drive's value always remains between 0 and 10. Figure 6 gives an illustration of how effects change the level of a drive.




Fig. 6. Impact of an action's effects on the level of a drive

The effects of drives are useful in the appraisal phase to create emotional impact that will be stored in the autobiographic memory and guide the agent's further actions. Since each agent has a different personality (refer to Section 4.2), the effect of an event may differ from one agent to another, which in turn affects their emotional and behavioural responses. For example, if the user performs an action that threatens the integrity of the Sprytes, character A (friendly) may choose to advise him/her kindly to not repeat the act while character B (hostile) may display anger towards the user. Since character A is friendly towards the user, the user may respond in a polite manner and perform actions that may restore character A's integrity, hence increases both its affiliation and integrity drives. As for character B, the user may choose to ignore it or perform more threatening behaviour that causes more harm. In this situation, character B may then perform actions that might reinstate its integrity such as apologising to the user. Thus, drives can be considered both the source of behaviour and a feedback from the effect of behaviour, a fundamental aspect necessary for learning agents.

Personality

By assigning different weights and decay factors for different drives to different ORIENT agents, characters with different personalities can be produced. For example, if character A is a friendly character, affiliation would be an important factor in its social relations, say weight 0.7 while a hostile character B would have a low importance for affiliation, say weight 0.3. What this means is that character A will have a higher tendency to show friendly behaviour compare to character B although responses may differ depending on circumstances and other factors such as like-dislike relationship. Now, if both characters have a current affiliation value of 6 and if an action with a negative effect on affiliation

(for instance -4) happens, character A's level of affiliation would be 3.2 ($6+(-4*0.7)$) while character B's level of affiliation would be 4.8 ($6+(-4*0.3)$) based on Equation 1. In this case, character A will work harder to satisfy its affiliation drive than character B. Additionally, a much higher decay factor on affiliation drive can be assigned to character A than to character B causing its affiliation strength to drop faster with time, further emphasising the importance of the drive to character A.

Although choosing the initial starting value for the parameters can be tricky, we have devised a mechanism to perform this based on psychological theories. First, the personality of our characters are defined by a group of psychologists using the BIG Five personality dimensions: neuroticism, extraversion, agreeableness, conscientiousness and openness [54]. Each dimensional parameter is assigned a value ranging from 1 to 5 for each character. For example, say character A above has value 1 for neuroticism, 4 for extraversion, 3 for agreeableness, 3 for conscientiousness and 4 for openness. These dimensions are then mapped onto drives. For example, the extraversion dimension is defined by high importance in affiliation drive while the conscientiousness dimension is defined by high importance in need for certainty and competence.

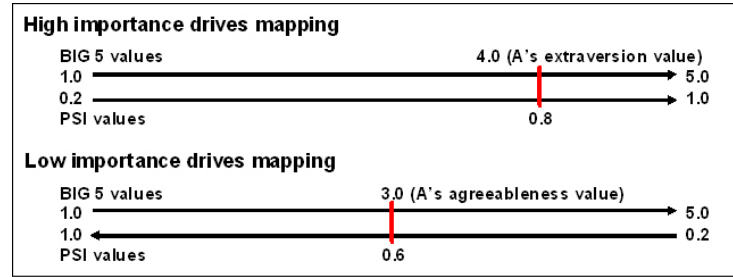


Fig. 7. Mappings from BIG Five to PSI for high and low importance drives

In order to distinguish between high importance and low importance drives, different mappings are adopted as shown in Figure 7. We assume that when drives have high importance, the value mapping for BIG Five to PSI is 1.0 - 5.0 to 0.2 - 1.0. When drives have low importance, the reverse mapping applies: 1.0 - 5.0 to 1.0 - 0.2. More than one dimension can affect a particular drive (e.g. the need for affiliation is affected by both the extraversion (high importance) and agreeableness (low importance) dimensions) and the final weight for each drive is calculated by averaging the mapped values for the relevant BIG Five dimensions. Thus, the weight for need of affiliation for character A would be $(0.8+0.6)/2=0.7$ (refer Figure 7). As a result, each character has a weight assigned to each need which defines the importance of the associated need to the character and influences its behavioural responses to different situations. This means different combination of weight values for needs will produce characters with distinguish-

ing behaviour, leading to a reflection of personality. More information about this mapping can be found in [55].

Cognitive drives: certainty and competence

As for certainty and competence, no explicit specification of contributions to these is necessary because they are cognitive drives and their values can be calculated automatically as described shortly. Whenever an expected event fails to turn up or an unknown object appears, the agent’s certainty drops. Thus, certainty represents the extent to which knowledge about a given fact/action is accurate or known. In order to model uncertainty about a given goal, ORIENT characters continuously make predictions about the probability of success of their goals (determined by the ratio of success and number of tries). These predictions are then compared with the actual outcomes (success or failure). The difference between these two values is the *ObservedError* (as illustrated in Figure 8).




Fig. 8. Modelling uncertainty

Using the current and past observed errors, we can determine the error or uncertainty in the prediction of goal success. Equation 2 describes how this is done using an Exponential Moving Average where the weighting factors decrease exponentially, resulting in the most recent data being the most important. α represents the rate at which past observations lose importance and t is the time step for the character’s mind cycle.

$$Uncertainty_t(g) = \alpha * ObservedError_t(g) + (1 - \alpha) * Uncertainty_{t-1}(g) \quad (2)$$

The Uncertainty formula is then used to determine the effect that a given goal has on the certainty drive according to equation 3. The calculation of this effect assumes that the goal will be successful, and as such $ObservedError_{t+1}$ is given by $1 - P(g)$, where 1 represents the predicted successful outcome and $P(g)$ is the goal’s success probability. Gaining certainty (and reducing uncertainty) corresponds to creating a more accurate model of the world. As such, goals that are relatively unexplored or that present higher error estimation (higher uncertainty) present more potential for improvement. Following this rationale, certainty is achieved by exploration of new strategies and actions, which leads to the construction of more complete hypotheses.

$$Effect(g, certainty) = Uncertainty_t(g) - Uncertainty_{t+1}(g) \quad (3)$$

Please note that the character does not learn by forming new goals because this will lead to a lack of control over its behaviour. Instead, it learns by trying out different actions from a pre-specified set of actions and remembering which actions helped it to tackle a situation best. For instance, in the previous example, character B will remember that showing anger towards the user will not help it to restore its integrity. This information is stored in its autobiographic memory and is used to determine the success probability of actions/goals in satisfying a specific drive in future. Since certainty depends on the amount of known information relating to a goal, the more an agent encounters the same type of situation, the higher its certainty is regarding the situation.




Fig. 9. Need for competence

Competence represents the efficiency of an agent in reaching its goals and fulfilling its demands. Success of goals increases competence while failure decreases it as depicted in Figure 9. There is an important simplification, which is to consider that all goals have the same contribution to competence, independently of their importance. So, the expected contribution of a given goal to competence depends exclusively on the goal's probability, as seen in Equation 4. The constant k defines the change in competence according to the goal's success or failure.

$$Effect(g, competence) = k \times P(g) - k \times (1 - P(g)) \quad (4)$$

A low competence level indicates that the agent should avoid taking risks and choose options that have worked well in the past (and thus have high probability) as shown in Figure 9. A high competence means that the agent can actively seek difficulties by experimenting with new courses of action that are less likely to succeed. Together, competence and certainty direct the agent towards explorative behavior; depending on its abilities and the difficulty of mastering the environment, it will actively seek novelty or avoid complexity.

Goal activation and Intention selection

At the start of an interaction, each agent has a set of initial values for drives. Based on the level of its current drives, the agent generates intentions, that is,

it activates goal(s) that are relevant to the perceived circumstances. In PSI only the likelihood of achieving a single drive and the level of that drive are taken into account when calculating a goal's utility. In ORIENT however, a drive may have several goals that satisfy it (e.g. a Spryte character can gain affiliation by befriending the user or socialising with another Spryte) and a goal can also affect more than one drive (e.g. eating food offered by another Spryte satisfies the need for energy as well as affiliation). So, when determining a goal's utility (Equation 5), all drives that it satisfies are taken into account. A goal that satisfies more drives will have a higher utility than those that satisfy fewer.

$$Utility(g) = \sum_d Utility(g, d) \quad (5)$$

The more a goal reduces a drive's deviation, the more important that goal is (e.g. eating satisfies an ORIENT character's need for energy more than drinking). The strength or relevance of a drive is dependent on the current situation. For example, the goal of eating will be very important if the character is extremely hungry and insignificant if the character has a high level of energy. Equation 6 depicts how the utility of a goal is determined for a given drive. It corresponds to the difference between the current relevance or strength of the drive ($Str(Level_t(d))$) and the expected relevance after the goal is achieved and the effect applied to the drive ($Str(Level_t(d) + effect(g, d))$). The strength of a drive is determined by the quadratic difference between the maximum value of the drive and its current level (as depicted in Equation 7). A non-linear function is applied to the utility function, $Utility(g, d)$, in order to ensure that the same effect of a goal has a higher relevance when the drive's level is low than when the drive's level is high. For instance, if a goal has an effect of +1, the utility value will be 1 if the current level of the drive is 9 and 17 if the current level of the drive is 1.

$$Utility(g, d) = Str(Level_t(d)) - Str(Level_t(d) + effect(g, d)) \quad (6)$$

$$Str(L) = (10 - L)^2 \quad (7)$$

The expected utility value of a goal ($EU(g)$) can then be determined by the utility the goal has on the character's drives and the expected competence in achieving the goal (as seen in Equation 8). Additionally, the urgency of a goal which is inversely proportional to the drive levels (0: not urgent or 1: very urgent) is taken into account. $GoalUrgency(g)$ gives importance to goals that should be activated immediately.

$$EU(g) = (1 + goalUrgency(g)) * ExpectedCompetence(g) * Utility(g) \quad (8)$$

A goal's expected competence represents the agent's likelihood in achieving the goal. A first approach could be to directly use the goal's success probability as a measure of expected competence (Equation 9). However, if by chance a goal

reached probability 0 (for instance by failing the first time it became active), it would never be selected again since the expected utility would always be 0. For this reason, the expected competence (Equation 11) of the agent will then be an average of its overall competence (determined by the ratio of success in achieving goals) and its competence in performing the particular goal. Thus, even if the goal's success probability is low, the agent can consider selecting it if it has a high general competence.

$$GoalCompetence(g) = P(g) = \frac{NoOfSuccesses(g)}{NoOfTries(g)} \quad (9)$$

$$OverallCompetence = \frac{NoOfSuccesses}{NoOfGoalsTried} \quad (10)$$

$$ExpectedCompetence(g) = \frac{OverallCompetence + GoalCompetence(g)}{2} \quad (11)$$

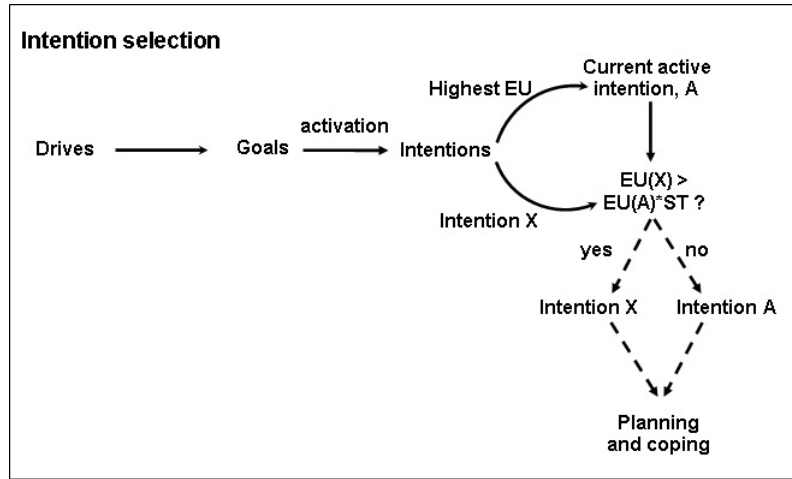


Fig. 10. Goal activation and intention selection

On each cycle, goals are checked to see if any has become active by testing the goal's preconditions. Once a goal becomes active, a new intention to achieve the goal is created and added to the intention structure. The intention represents the agent's commitment to achieve the goal and stores all plans created for it. Since there can be more than one intention activated at any particular time instance, the character must choose one of them to continue deliberation (and planning). Applying PSI, the selection of goals in ORIENT is performed based on the *selection threshold* value. The current active intention is selected based on a winner takes all approach, that is, the goal with the highest expected utility value (according to Equation 8) is chosen. An unselected goal can be activated if

its expected utility surpasses the value of the current active intention multiplied by the *selection threshold*. Figure 10 gives an illustration of this selection process. So, if the *selection threshold* is high, it is less likely that another goal will be activated, hence, allowing the agent to concentrate on its current active intention.

Planning and coping

After an intention is selected, the agent proceeds to generate plan(s) to achieve it. When a plan is brought into consideration by the reasoning process, it generates and updates OCC prospect based emotions such as:

- Hope: 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: Fear of not being able to achieve the intention. The emotion intensity is determined from the goal’s importance of failure and the plan’s probability of failing.

As already mentioned, FATiMA employs a continuous planner which monitors all events continuously in order to detect when an action is accomplished or fails. It also handles unexpected events that affect future plans, and updates all plans accordingly (e.g. when some preconditions for a goal have been achieved due to actions of another agent). This process will change the character’s internal plan as well as plan probabilities leading to different emotional appraisals of Hope and Fear.

A continual check is also performed on all active goals to determine if the goal succeeds or fails. If the planner is unable to make a plan, more prospect based emotions will be generated, such as Satisfaction, Disappointment, Relief and Fears-Confirmed. In order to cope with different circumstances, ORIENT characters perform two types of coping: problem-focused coping and emotion-focused coping as in FATiMA. These coping strategies are again triggered by emotions and personality of the characters. For instance, a fearful Spryte has a higher chance of dropping an uncertain goal than a hopeful character.

5 First Prototype

A prototype version of the ORIENT software has been implemented and tested. There are three scenarios in the game - meal, recycling and gardening. In each scenario, the player (in the game, all three users are assumed to be one because they co-operate to submit one input at a time into the system) has a specific task to achieve and his/her main aim is to interact with the Sprytes and adapt him/herself to the Sprytes culture so that collaboration to save the planet can be established. During interaction with the user, a Spryte character’s drives change continuously depending on its own actions and the user’s. Figure 11 shows a short interaction screen shot between the user and one of the characters, Abbuk in the meal scenario.



Fig. 11. A short interaction between the user and Abbuk

Figure 12 presents Abbuk's drives at the start of the simulation while Figure 13 shows Abbuk's drives after the simulation ended. It can be observed that by the end of simulation, all of Abbuk's drives except certainty have increased in values. This is because the greeting goal contributes positively to the affiliation drive. From time to time, Abbuk rests to maintain/increase its energy and integrity level. And since there is no threatening situation, Abbuk's competence level also increases due to its ability to cope with the situation. However, its certainty level has dropped slightly due to decay and the non-performance of highly uncertain goal that may boost certainty.

Once Abbuk starts interacting with the user, it also makes prediction about the user's drive values so that it can take appropriate action considering its cultural attributes [22]. For example, if it is from a collectivist culture, then, the user's needs are considered as important as its own [23]. Therefore, in Figure 13, the user's expected drive values are listed. In this simulation, a simple user interface that consists of a drop down box for user action options is used. In the real game, users interact with the character using a mobile phone (speech input and RFID tag scanning), a Dance Mat (navigation) and a WiiMote (gesture input). For more information, please refer to [24, 21].

Although no evaluation has yet been carried out to test the technical improvement of this architecture over FAtiMA or PSI in terms of authoring or the characters' behaviour, - difficult to do in isolation - an evaluation of the game, ORIENT was performed. Briefly, the evaluation of ORIENT was designed as an in-role experience for adolescents in UK and Germany focusing on two inter-related themes: participants' intercultural awareness; and their immersion and engagement in interactions with ORIENT. Four groups of three adolescents in each country took part in sessions of approximately two hours. Rather than being presented as explicit evaluation, the instruments used were incorporated

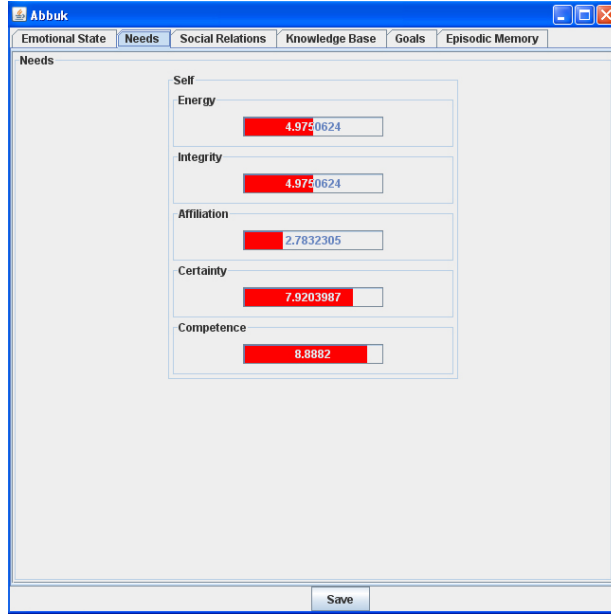


Fig. 12. Abbuk’s drive values before interacting with the user

into the role-play as Space Command debriefing and as support for the users mission to gain greater awareness of the Spryte culture. This approach allows data capture in an interesting and engaging way. For instance, participants are required to fill an applicant form (demographic data) and the Cultural Intelligence Scale to show that they are “qualified” for the internship. More details on the instruments employed and the experimental setting can be found in [21, 56].

In post-interaction discussions, the participants reported that they found the Sprytes and their culture engaging. They were able to identify similarities and differences between their own and the culture of the Sprytes, suggesting the basis of empathy exists. They added that these differences appeared to be believable and credible which is consistent with our aim to create adaptive affective autonomous character. Moreover, they are very interested in participating in the Sprytes’ activities. Despite the positive feedback, it was found that believability as well as the emotional impact of the drama on the users could be further fostered by giving the individual Sprytes more distinctive personalities as they currently appear as cultural stereotypes.

6 Conclusion

In this paper, we discussed work in developing adaptive affective autonomous NPCs for an educational role play game. Of crucial importance in the system described is the integration of a motivational system and an appraisal system

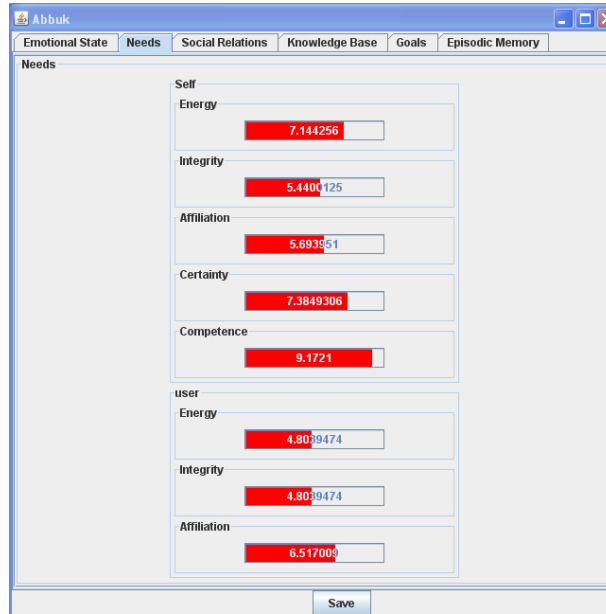


Fig. 13. Abbuk's drive values after interacting with the user

with the planning mechanism. In ORIENT, characters behaviour is regulated by a biologically-inspired architecture of human action regulation combining FATiMA and PSI so that we can take advantage of the benefits of both systems while avoiding the flaws. The resulting architecture is flexible and robust since characters are now able to learn and adapt their behaviour to different environmental circumstances depending on their needs. Moreover, the emergence emotions resulting from and leading to various behavioural responses generates a whole range of emotional and empathic outcomes. The new addition of the motivational system to FATiMA provides ORIENT characters with a basis for selective attention where all behaviour ultimately serves the purpose of achieving the character's drives which can be physiological, social or cognitive. Intentions are selected based on strength of drives, urgency and success probability addressing the BDI architecture's lack of reasoning in why certain decisions are chosen while others are discarded. The resulting characters learn through trial and error, allowing more efficient adaptation and empathic engagement in different social circumstances. The successful linking of body and mind is consistent with that of humans' and hence, should produce characters with behaviours that seem plausible to a human. This is proven to some extent in the user evaluation of the game. Moreover, the designers are relieved from extensive authoring tasks. The software, which is written in Java, has been made available at the open source portal (SourceForge⁸ and is reusable in autonomous agents applications.

⁸ <http://sourceforge.net/projects/orient-ecircus>

Acknowledgements

This work was partially supported by a scholarship (SFRH BD/19481/2004) granted by the Fundação para a Ciência e a Tecnologia (FCT) and by European Community (EC) and was funded by the eCIRCUS project IST-4-027656-STP with university partners Heriot-Watt, Hertfordshire, Sunderland, Warwick, Bamberg, Augsburg, Wuerzburg plus INESC-ID and Interagens. The authors are solely responsible for the content of this publication. It does not represent the opinion of the EC or the FCT, which are not responsible for any use that might be made of data appearing therein.

References

- [1] Woodcock, S.: Game AI: The state of the industry. *Game Developer* **7**(8) (2000) 24–32
- [2] Castle, L.: The making of blade runner, soup to nuts! In: *Proceedings of the Game Developer’s Conference (GDC)*, Long Beach, CA (1998) 87–97
- [3] Orkin, J.: Three states and a plan: The a.i. of f.e.a.r. In: *Proceedings of the Game Developer’s Conference (GDC)*, San Jose (2006) Available at http://web.media.mit.edu/~jorkin/gdc2006_orkin_jeff.fear.pdf, Last accessed: Aug 13, 2009.
- [4] Bates, J.: The role of emotion in believable agents. *Communications of the ACM: Special Issue on Agents* **37**(7) (Jul 1994) 122–125
- [5] Picard, R.W.: *Affective Computing*. MIT Press (1997)
- [6] Thomas, F., Johnston, O.: *Disney Animation: The Illusion of Life*. New York: Abbeville Press (1981)
- [7] Minsky, M.: *Society of Mind*. Simon and Schuster (1988)
- [8] Damasio, A.: *Descartes’ Error: Emotion, Reason and the Human Brain*. Gosset/Putnam Press, New York (1994)
- [9] Moffat, D.: Personality parameters and programs. In Petta, P., Trappl, R., eds.: *Creating Personalities for Synthetic Actors, Towards Autonomous Personality Agents*. Springer-Verlag (1997) 120–165
- [10] Jones, C.: *Chuck Amuck: The Life and Times of an Animated Cartoonist*. Farrar, Straus and Giroux (1989)
- [11] Paiva, A., Dias, J., Aylett, R., Woods, S., Hall, L., Zoll, C.: Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence* **19** (2005) 235–266
- [12] Core, M., Traum, D., Lane, H., Swartout, W., Gratch, J., Lent, M.v., Marsella, S.: Teaching negotiation skills through practice and reflection with virtual humans. *Simulation* **82**(11) (2006) 685–701
- [13] Tomlinson, B., Blumberg, B.: Alphawolf: Social learning, emotion and development in autonomous virtual agents. *First GSFC/JPL Workshop on Radical Agent Concepts* (Oct, 04 2002)
- [14] Bickmore, T., Picard, R.: Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction* **12** (2005) 293–327
- [15] Kenny, P., Parsons, T., Gratch, J., Leuski, A., , Rizzo, A.: Virtual patients for clinical therapist skills training. In Pelachaud, C., Martin, J., André, E., Chollet, G., Karpouzis, K., Péle, D., eds.: *7th International Conference on Intelligent Virtual Agents*. *Lecture Notes in Computer Science*, Paris, Springer (2007) 192–210

- [16] Eck, R.V.: Digital game-based learning: It's not just the digital natives who are restless. *EDUCAUSE Review* **41**(2) (2006) 16–30
- [17] Piaget, J.: *The Principles of Genetic Epistemology*. Routledge and Keegan Paul Ltd, London (1972)
- [18] Lewin, K.: *Field Theory in Social Science*. Harper and Row, New York (1951)
- [19] Kanfer, F.H.: Self-management strategies and tactics. (1979) 195–224
- [20] Dias, J., Paiva, A., Vala, M., Aylett, R., Woods, S., Zoll, C.: Empathic characters in computer-based personal and social education. In Pivec, M., ed.: *Affective and Emotional Aspects of Human-Computer Interaction Game-Based and Innovative Learning Approaches*. IOS Press (2006)
- [21] Lim, M.Y., Aylett, R., Enz, S., Kriegel, M., Vannini, N., Hall, L., Jones, S.: Towards intelligent computer assisted educational role-play. In Chang, M., Kuo, R., Kinshuk, Chen, G.D., Hirose, M., eds.: *4th International Conference on E-Learning and Games, Edutainment 2009, Banff, Canada, Springer-Verlag Berlin Heidelberg 2009* (2009) 208–219
- [22] Mascarenhas, S., Dias, J., Afonso, N., Enz, S., Paiva, A.: Using rituals to express cultural differences in synthetic characters. In Decker, Sichman, Sierra, Castelfranchi, eds.: *8th International Conference on Autonomous Agents and Multiagent Systems, Budapest, Hungary, International Foundation for Autonomous Agents and Multiagent Systems* (2009) 305–312
- [23] Hofstede, G.: *Cultures and Organisations*. McGraw-Hill, London (1991)
- [24] Lim, M.Y., Kriegel, M., Aylett, R., Enz, S., Vannini, N., Hall, L., Rizzo, P., Leichtenstern, K.: Technology-enhanced role-play for intercultural learning contexts. In Natkin, S., Dupire, J., eds.: *8th International Conference on Entertainment Computing, Paris, France, IFIP International Federation for Information Processing 2009* (2009) 73–84
- [25] Rime, B.: Interpersonal emotion regulation. In Gross, J., ed.: *Handbook of Emotion Regulation*, Guilford Press (2007)
- [26] Hogan, R.: Development of an empathy scale. *Journal of Consulting and Clinical Psychology* (35) (1977) 307–316
- [27] Hoffman, M.L.: Empathy, its development and prosocial implications. *Nebraska Symposium on Motivation* **25** (1977) 169–217
- [28] Ho, W.C., Dautenhahn, K., Nehaniv, C.L.: Computational memory architectures for autobiographic agents interacting in a complex virtual environment: A working model. *Connection Science* **20**(1) (2008) 21–65
- [29] Oliveira, E., Sarmento, L.: Emotional advantage for adaptability and autonomy. In: *AAMAS, Melbourne, Australia, ACM* (2003) 305–312
- [30] Cañamero, D.: A hormonal model of emotions for behavior control. In: *VUB AI-Lab Memo 97-06, Vrije Universiteit Brussel, Belgium* (1997)
- [31] Velásquez, J.D.: Modeling emotions and other motivations in synthetic agents. In: *Proceeding AAAI 97, AAAI Press and The MIT Press* (1997) 10–15
- [32] Blumberg, B.: *Old Tricks, New Dogs: Ethology and Interactive Creatures*. PhD thesis, Massachusetts Institute of Technology, MIT, Cambridge, MA (1996)
- [33] Pezzulo, G., Calvi, G.: Rational agents that blush. In Paiva, A., Prada, R., Picard, R., eds.: *Proceedings of the 2nd International Conference of Affective Computing and Intelligent Interaction. Lecture Notes in Computer Science, Springer* (2007) 374–385
- [34] Izard, C.E.: Four systems for emotion activation: Cognitive and noncognitive processes. *Psychological Review* **100**(1) (1993) 68–90
- [35] Arnold, M.B.: *Emotion and Personality. Volume 1*. Columbia University Press (1960)

- [36] Ortony, A., Clore, G., Collins, A.: The cognitive structure of emotions. Cambridge University Press, Cambridge, UK (1988)
- [37] Elliot, C.D.: The Affective Reasoner: A process model of emotions in an multi-agent system. PhD thesis, Northwestern University, Illinois (1992)
- [38] Reilly, W.S., Bates, J.: Building emotional agents. Technical Report CMU-CS-91-143, School of Computer Science, Carnegie Mellon University (1992)
- [39] Loyall, A.B., Bates, J.: Hap: A reactive adaptive architecture for agents. Technical Report CMU-CS-91-147, School of Computer Science, Carnegie Mellon University (1991)
- [40] Gratch, J., Marsella, S.: Evaluating a computational model of emotion. *Journal of Autonomous Agents and Multiagent Systems* (Special issue on the best of AAMAS 2004) **11**(1) (2004) 23–43
- [41] Dias, J., Paiva, A.: Feeling and reasoning: A computational model for emotional agents. In: 12th Portuguese Conference on Artificial Intelligence (EPIA 2005), Portugal, Springer (2005) 127–140
- [42] Dörner, D.: The mathematics of emotions. In Frank Detje, D.D., Schaub, H., eds.: *Proceedings of the Fifth International Conference on Cognitive Modeling*, Bamberg, Germany (2003) 75–79
- [43] Dörner, D., Gerdes, J., Mayer, M., Misra, S.: A simulation of cognitive and emotional effects of overcrowding. In Fum, D., Del Missier, F., Stocco, A., eds.: *Proceedings of the 7th International Conference on Cognitive Modeling*, Trieste, Italy, Edizioni Goliardiche (2006)
- [44] Sloman, A.: Varieties of affect and the cogaff architecture schema. In: *Symposium on Emotion, Cognition and Affective Computing*, AISB 01 Convention, University of York, United Kingdom (Mar, 21-24 2001) 39–48
- [45] Randolph M. Jones, Amy E. Henninger, E.C.: Interfacing emotional behavior moderators with intelligent synthetic forces. In: *Proceeding of the 11th CGF-BR Conference*, Orlando, FL (May, 7 2002)
- [46] Bratman, M.E.: *Intention, Plans and Practical Reasoning*. Harvard University Press, Cambridge, Massachusetts (1987)
- [47] Aylett, R., Dias, J., Paiva, A.: An affectively driven planner for synthetic characters. In: *International Conference on Automated Planning and Scheduling (ICAPS2006)*, UK, AAAI Press (2006) 2–10
- [48] Marsella, S., Johnson, B., LaBore, C.: Interactive pedagogical drama. In: *Fourth International Conference on Autonomous Agents (AAMAS)*, Bologna, Italy, ACM Press (2002) 301–308
- [49] Dias, J., Ho, W.C., Vogt, T., Beeckman, N., Paiva, A., André, E.: I know what I did last summer: Autobiographic memory in synthetic characters. In Paiva, A., Prada, R., Picard, R., eds.: *Affective Computing and Intelligent Interaction*, Lisbon, Portugal, Springer (2007) 606–617
- [50] Aylett, R., Vannini, N., André, E., Paiva, A., Enz, S., Hall, L.: But that was in another country: agents and intercultural empathy. In Decker, Sichman, Sierra, Castelfranchi, eds.: *8th International Conference on Autonomous Agents and Multiagent Systems*, Budapest, Hungary, International Foundation for Autonomous Agents and Multiagent Systems (2009) 305–312
- [51] Bartl, C., Dörner, D.: Comparing the behavior of psi with human behavior in the biolab game. In Ritter, F.E., Young, R.M., eds.: *Proceedings of the Second International Conference on Cognitive Modeling*, Nottingham, Nottingham University Press (April 1–4 1998)

- [52] Lim, M.Y.: Emotions, Behaviour and Belief Regulation in An Intelligent Guide with Attitude. PhD thesis, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, Edinburgh (2007)
- [53] Squire, K.: Design principles of next generation digital gaming for education. *Educational Technology* **43**(5) (2003) 17–33 The Games-To-Teach Team at MIT.
- [54] McCrae, R.R., John, O.P.: An introduction to the five-factor model and its applications. *Journal of Personality* **60** (1992) 175–215
- [55] Nazir, A., Enz, S., Lim, M.Y., Aylett, R., Cawsey, A.: Culture-personality based affective model. *AI and Society Journal* (2009) Available at: <http://www.springerlink.com/content/4x35613r5k55q1q5/>.
- [56] Enz, S., Zoll, C., Vannini, N., Lim, M.Y., Hall, L., Paiva, A., Aylett, R.: Orient - the intercultural empathy through virtual role-play. In Mukerji, S., Tripathi, P., eds.: *Cases on Transnational Learning and Technologically Enabled Environments*. IGI Global (2009) 65–88