I Want to Be Your Friend: Establishing Relations with Emotionally Intelligent Agents

João Dias, Ana Paiva
Universidade Técnica de Lisboa, Instituto Superior Técnico, INESC-ID
Taguspark, Av. Prof. Cavaco Silva
2780-990 Porto Salvo, Portugal
{joao.dias,ana.paiva}@inesc-id.pt

ABSTRACT

This paper addresses the problem of how to create autonomous virtual agents that are able to intentionally establish and strengthen social relations with other agents and humans. To do so, a computational model is proposed that embraces explicit intentions in the agents that are aimed at increasing the other's relations towards them. The model is inspired in the notion of emotional intelligence, and allows for agents to reason about the emotions of others and perform interpersonal emotion regulation (IER) in order to dynamically create the relations with others. The model was used to build the NPCs of a computer game scenario in NeverWinter Nights 2, where the player is required to work together with two agents to achieve a particular quest. An experiment was then conducted where players interacted with either a version with or without IER. By measuring friendship quality we were able to show that when one of the agents employed IER strategies it was perceived as more friendly according to four out of six dimensions of friendship.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent Agents*; J.4 [**Social And Behavioral Sciences**]: Sociology

Keywords

Virtual Agents; Emotional Intelligence; Emotion Regulation; Relational Behaviour

1. INTRODUCTION

Intelligent Virtual Environments are becoming everyday more widespread in our society covering many different application areas. Examples can be found for intelligent learning environments [16] or interactive storytelling systems [3]. Many of these environments embrace the presence of intelligent virtual agents (IVAs), that "autonomously" act in an intelligent and believable manner, interacting not only with other agents but also with the users. Yet, as the complexity of these environments grows, the expected "intelligence" of such virtual agents also needs to grow. As such, it is important for intelligent virtual agents to be rich in the portraying of social and emotional behaviours. To that end, they need to embed features that will allow them to automatically and autonomously establish social relations with one another or even with the user. This is particularly important in storytelling systems and role-playing games since the existence and development of social relations is fundamental to create engaging narratives.

Taking this into consideration, this paper addresses the problem of creating autonomous virtual agents that are able to intentionally establish and strengthen social relations with other agents and humans. Although it is possible for users to create social relations with simple agents, in this research we are interested in modeling agents that use explicit intentions and actions to increase the other's relations towards them. Furthermore, the model proposed here aims at being flexible and generic enough to allow us to easily deploy such agents in different scenarios where different actions are available.

To achieve this goal, we need to explore mechanisms that mimic the way that relations between real people evolve. To do this, we draw inspiration from the work of Salovey and Mayer [23], that define emotional intelligence as the understanding of the impact that emotions have on the self and on others, how emotions are created, and also being able to use this knowledge to regulate emotions on the self and in others. But how is emotional intelligence related to the development of social relations? Several studies [11, 10] have shown that people who have higher emotional intelligence have more positive social interactions with peers, and students with better scores in managing emotions in others were more liked and valued by the opposite sex[11]. Thus, it seems that the high emotionally intelligent individual is rather successful at establishing social relations with others. Therefore, if we are able to understand how regulating emotions in others affects relations, and if we are able to model this capability in our agents we can hope to have better results at dynamically establishing social relations in our synthetic virtual agents.

To that end we propose an emotionally intelligent agent model, which has an explicit model of Social Relations, is able to reason about emotions of others and perform interpersonal emotion regulation in order to dynamically create relations with others. The model was used to build a scenario in NeverWinter Nights 2, where the player is required to work together with two autonomous agents to achieve a particular quest. An experiment was then conducted by having players interacting either with a version with emotion regulation or without, and measuring friendship quality.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

2. ON EMOTIONAL INTELLIGENCE AND EMOTION REGULATION

In order to understand what are the requirements needed to endow agents with the capability to reason about emotions and establish social relations, it is necessary to take a deeper look at concepts such as Emotional Intelligence and Emotion Regulation. Mayer and Salovey [9] put forward a four branch model that further divides Emotional Intelligence into four main skills: 1) perceiving emotions in oneself and others - has to do with the perception and expression of emotion through gestures and facial expressions; 2) using emotions to facilitate thought - that focuses on using emotions to guide cognitive processes, such as learning and decision making; 3) understanding emotions - involves understanding the meaning of emotions and reasoning about them; and 4) regulating emotions - managing one's own and other's emotions in order to promote social goals (e.g. doing something pleasant to help a friend to overcome some bad mood). All but the second skill are directly relevant to our goal, interpersonal emotion regulation has an impact on the development of social relations, and it cannot be achieved in a automatic way without endowing our agents with the capability to understand and reason about emotions, and with the capability of knowing and representing other's emotional states. Therefore, these three skills need to be addressed in our proposed model.

Gross further divides emotion regulation into five families of processes (as seen in Fig. 1), according to the point in time where they affect the emotion generation process (appraisal):

- Situation selection - is the earliest process and corresponds to taking actions that make it more/less likely to end up in a situation we expect to give rise to a desirable/undesirable event.
- Situation modification - once the situation has occurred, situation modification works by externally modifying the situation so as to alter its emotional impact. For instance, by removing a threat that is causing a distress emotion.
- Attentional deployment - two types of attentional deployment strategies are distraction and concentration. Distraction focuses attention on different aspects of the situation or moves attention away from the whole situation, while concentration draws attention to the emotional features of a situation. The idea here is that if the situation (or relevant features of the situation) is not under attention, it will not be appraised, thus not generating the corresponding emotions.
- Cognitive change - works by internally changing the interpretation one has about the event or situation, so that it changes the resulting emotional outcome. It may involve reappraising the event (thinking to one self that what happened is not so bad after all), trying to find positive aspects in the situation, and other emotion-focused coping strategies such as the ones proposed by Marsella in [13].
- Response modulation - occurs after response tendencies have been initiated. It corresponds for instance, to one trying to regulate expressive behaviour when being happy in a situation where that is not socially acceptable to express such happiness.

Following a different approach, Niven, Totterdell and Holman [17] tried to determine what kind of strategies are used in everyday life by people to regulate affective processes in others. To that end they have set up a set of questionnaires and obtained 378 (e.g. validation, compliment) different

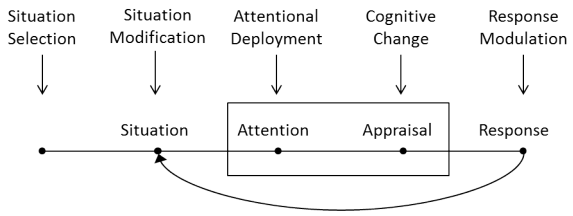

Figure 1: Five families of emotion regulation strategies

strategies identified by people. These strategies were then grouped into two main categories: engagement strategies which focus on the problem that caused the emotion, and relationship-oriented strategies. According to [6, 24], collateral effects of emotional episodes that arise from unattainment of goals, result in an increase of socioaffective needs such as social support and reassurance. In these situations, relationship-oriented strategies can be used to increase a target's socioaffective needs.

3. RELATED WORK

The problem of creating social relations in autonomous agents has already been addressed in many different ways. For instance Psychsim [22] is an environment for multi-agent simulation that employs a formal decision-theoretic approach using recursive models. This allows the agents to reason about the behavior and beliefs of other agents. Psychsim was used to model social-simulation scenarios and is able to create behaviour in order to attain a desired social relation state, however it does not have an explicit model of emotions. The SGD Model proposed by Prada [21] focus on modeling group dynamics between social agents. Amongst other things, it models relations of social power and of social attraction which are used to help build the dynamic of the group and resulting group behaviour. SGD model uses Heider's balance theory to determine that when an agent observes a positive socio-emotional interaction towards him, then his attraction for the performer of the interaction will increase. Ochs et al. [18] have proposed a model for NPCs for computer games, where the intensities and the types of emotions triggered by appraisals are used to update each of the social relations. For instance, a positive emotion caused in i by j increases the degree of liking that i has for j while a negative emotion decreases it. Unfortunately, Ochs et al. model does not address how social relations and emotions should be used to affect behaviour.

There are currently a few systems or models that perform emotion regulation. Adam [1] presents a formalization of a cognitive logical model of emotions, which allows us to reason about, and plan for an agent's emotions. It also proposes a formalization of coping behaviour, which can be used to regulate emotions in others. However, there is no explicit model for social relations (social behaviour is implicit in the rules defined). Boss and Lange [4] aim at developing IVAs with a Theory of Emotion Regulation (ToER), i.e, a Theory of Mind that models emotional states of others, and which allows them to reason about other's emotion regulation processes. Although the ToER model uses explicit interpersonal emotion regulation strategies, it does not model social relations and does not have explicit social goals. EmoEmma

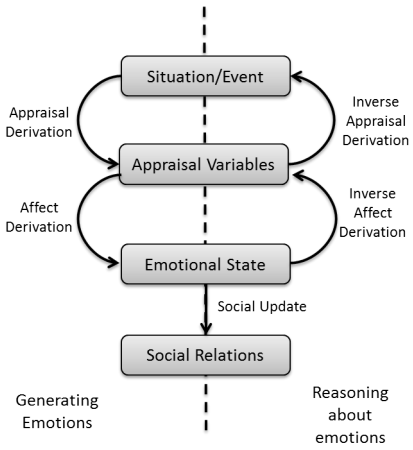

Figure 2: A model for Emotionally Intelligent Agents that reasons about and regulates emotions

[20] is an Interactive Storytelling System where the characters are driven by higher-level goals that correspond to desired states of mind (e.g. seeking happiness) formalized as conjunction of feelings. Unfortunately, since there is no appraisal process explicitly modeled in EmoEmma, we need to specify all actions that can generate a particular emotion. This mechanism is not flexible, for instance not allowing us to have a context-dependent subjective appraisal process.

4. AN EMOTIONALLY INTELLIGENT AGENT MODEL

As described previously, we aim at creating agents that dynamically establish and manage social relations with other agents. In our approach we are endowing agents with three of the emotional intelligence skills identified by Salovey and Mayer. To that end, our model needs to: generate and express emotions; reason about emotions; and regulate emotions in others (i.e. perform interpersonal emotion regulation). Furthermore, to properly address the establishment of social relations with others it is also necessary to model the existing relations between agents, and the influence of emotions in their development.

Figure 2 shows the proposed model. It presents a series of simplifications regarding Gross's model of emotion regulation families. Although interesting, in practice performing Cognitive Change or Response Modulation in others is not an easy task. As Adam pointed out, it requires the existence of a large pool common-sense knowledge and the ability to infer positive aspects of a situation [1]. And this is not the focus of our work. For this reason, the proposed model focuses solely on families of regulation where the agent has more control: Situation Selection and Situation Modification. Since these families work in a similar way by performing actions to change the current situation, they where grouped together in our model.

The processes depicted in the model can be executed in both directions. When performed top-bottom, it corresponds to the appraisal process which receives an event and generates emotions, while when performed bottom-top it allows the agent to reason about the appraisal process, allowing it to determine what events can cause a particular desired emotion in a target. The last stage of the model represents the connection between emotions experienced and the development of social relations.

Note that in order for model to reason about the emotional state of another agent, it needs to have a model of that agent's emotional state in the first place. This is accomplished by a Theory of Mind Mechanism that represents, stores and updates information about other's internal states (such as emotions and social relations) using a recursive version of the model proposed. A complete description of the Theory of Mind Mechanism is out of the scope of this paper and can be found in [5]. We will now describe the several processes in more detail.

Generating Emotions: the Appraisal Process

The appraisal mechanism proposed derives from appraisal theories, where emotions (or emotional responses) are seen as valenced (i.e. good or bad) reactions to events that arise from a subjective appraisal process. However, it is important to point out that the model proposed does not commit itself with any particular appraisal theory. The appraisal mechanism was divided into two separate processes, following the distinction proposed by [12]. The first one, appraisal derivation, is responsible for evaluating the relevance of the event to the agent well being and determines a set of appraisal variables (for instance, OCC Appraisal Theory [19] models variables such as desirability and desirability for others). The second process, affect derivation, takes the appraisal variables as input and generates the resulting emotions (and corresponding intensities) according to a particular appraisal theory. As example, in the case of OCC Theory of emotions an event appraised as desirable for self and undesirable for other will generate an emotion of type Gloating. This explicit division into distinct processes of appraisal facilitates the process of reasoning about emotions by allowing the agent to reason about each part of the appraisal process separately.

Modeling Social Relations

In order for our agents to be able to create social relations, we need an explicit model of social relations. Here we model the relation of social attraction, first studied by Moreno [15], which reflects the affective ties (or degree of liking) that one person (or agent in our case) establishes with the others. Attraction, or liking towards a target t is represented by a number that can be positive (representing liking) or negative (representing dislike). Attraction is not necessarily reciprocal, agent A may like B, but agent B may dislike agent A.

The dynamic of social relations (i.e. the way how relations evolve according to interactions) follows from Heider's Balance theory [8]. Balance Theory hypothesis is that people avoid inconsistent cognitive configurations and that they mobilize their efforts to change it to a consistent state. As example, if someone you are attracted to likes a particular thing you dislike, you will either tend to start liking that thing more (or dislike less), or loose some of the attraction to that person.

This same principle is applied to the perception of events, similarly to SGD model. Suppose that the user performs an action undesirable for the agent, such as insulting him. Since the event is performed by the user, then we assume that the event is desirable to him. If the agent likes the

user, this corresponds to an inconsistent state. Differently from SGD model, events are not pre-classified as positive or negative. Instead, we use the result of the appraisal to determine whether a perceived event has a positive or negative value, in a similar way to Ochs et al. model. This value is then used to determine the change in the current existing relation. More specifically, whenever a new emotion *em* caused by another agent is added to the emotional state, the social relation towards that agent is updated according to equation 1, in order to converge to a more stable state.

$$Like(target)_{t+1} = Like(target)_t + k \times V(em) \times I(em) \quad (1)$$

V(em) is a function that returns the valence of the emotion *em* (+1 for positive emotions, -1 for negative emotions), while I(em) returns the intensity of the emotion. k is a constant ranged between 0 and 1, representing how strong should be the shift towards a consistent state. By looking at the formula, it is straightforward to acknowledge that positive emotions increase the interpersonal relation, while negative emotions decrease it. Stronger emotions cause bigger changes in the relationship.

To determine other's degree of liking towards the agent, the Theory of Mind Mechanism simulates the other's appraisals of actions performed by the agent, and uses such appraisals to update the relations as described.

Reasoning about Emotions

The model proposed was integrated on top of a continuous planner that builds plans of actions (and executes them) to achieve a desired goal state [2]. As such, in order to allow the agent to reason about emotions, it is necessary to represent the appraisal processes into STRIPS-like operators [7]. Two special Emotional Meta-Operators were created to represent the processes of appraisal derivation and affect derivation, as shown in Figure 3. The AppraisalDerivation Meta-Operator has the precondition that a particular event has happened (with subject $[s]$, action $[a]$ and target $[t]$) and was perceived by agent $[ag]$, and the effects that the event generated the appraisal variables $[var_1], .., [var_n]$ with values $[value_1], ..., [value_n]$ in agent $[ag]$. The AffectDerivation Meta-Operator specifies as preconditions that the agent currently has a set of appraisal variables, and as effect that a particular emotion $[emo]$ with valence $[valence]$ and intensity $[i]$ becomes part of the emotional state of the agent.

Unlike normal operators, Meta-Operators can bee seen as a template for operators that when selected can generate several instantiated operators. This is done by assigning an internal function to each of the Appraisal Meta-Operators. This internal function corresponds to inverting the appraisal processes in order to be able to reason backwards about emotions. Thus, whenever a Meta-Operator is instantiated in order to achieve a condition, it will use the corresponding inverse function to generate the concrete values for its own preconditions depending on its effects. The specific functions used will depend on the Appraisal Theory used to implement the appraisal processes.

To help the reader better understand this reasoning process, we will now provide a complete example of planning using the Emotional Meta-Operators. Suppose that the agent has the goal to make another agent B Joyful with intensity 3. The planner starts by selecting the AffectDerivation Meta-Operator and instantiates it with the substitutions

OPERATOR:
[ag]:AppraisalDerivation()
PRECONDITIONS:
[ag]:PerceivesEvent([s],[a],[t])
EFFECTS:
[ag]:Appraisal([var_1],[$value_1$])
...
[ag]:Appraisal([var_n],[$value_n$])

(a) Appraisal Derivation Operator

OPERATOR:
[ag]:AffectDerivation()
PRECONDITIONS:
[ag]:Appraisal([var_1],[$value_1$])
...
[ag]:Appraisal([var_n],[$value_n$])
EFFECTS:
[ag]:EmotionalState([emo],[valence],[i])

(b) Affect Derivation Operator

Figure 3: Representing Emotional Processes as Meta-Operators

$\{[ag]/B, [emo]/Joy, [valence]/positive, [i]/3\}$. Once instantiated, the operator automatically uses the Inverse Affect Derivation function to determine the appraisal variables and corresponding values necessary to create a Joy emotion. Assuming that we are using the OCC model, then the operator generated has the precondition B:Appraisal(Desirability, 3), meaning that agent B must have an appraisal with Desirability 3. This instantiated operator is added to the plan. Planning continues, and the planner selects the AppraisalDerivation MetaOperator to achieve the Joy Affect Derivation's precondition. Once more, there is an instantiation process by using the Inverse Appraisal Derivation function. The resulting operator has an effect of desirability 3, and has a precondition that a particular event e (e.g. give medicine to the target) considered desirable by agent B with value 3 was perceived by agent B. To finish up the plan, the planner will add to the plan an action to give the potion to agent B in order to achieve the AppraisalDerivation operator's precondition.

It is important to notice that the instantiation process of the AppraisalDerivation and AffectDerivation operators can generate more than one alternative concrete instantiation. In the above example, there may be more than one event that is appraised as having a desirability of 3. In this situation, the several different instantiations of the operator are added to different alternative plans, representing the alternative ways of achieving the desired appraisal variable.

Performing Interpersonal Emotion Regulation

The last issue to address is to decide when and how to perform interpersonal emotion regulation strategies. This is achieved by modeling explicit goals to establish or maintain social relations, which will trigger the activation of interpersonal emotion regulation, and the selection of the most appropriate strategies. Goals are defined with a set of preconditions that determine when the goal becomes active and a set of goal state conditions. There are currently two main Interpersonal Emotion Regulation goals used, depicted in figure 4: Proactive Positive Regulation and Reactive Positive Regulation.

In the Proactive Positive Regulation, the agent takes the initiative to regulate emotions in others (even if they do not necessarily feel bad), in order to make others like him.

GOAL:
ProactivePositiveRegulation([target])
PRECONDITIONS:
Like([target])>0
[deltaL]=Like([target])-[target]:Like(SELF)
[deltaL]>0
GOALSTATE:
[target]:EmotionalState([emo],positive,[deltaL])

(a) Proactive Positive Regulation Goal

GOAL:
ReactivePositiveRegulation([target])
PRECONDITIONS:
Like([target])>0
[target]:EmotionalState($[em_1]$,negative,[i])
GOALSTATE:
[target]:EmotionalState($[em_2]$,positive,[i])

(b) Reactive Positive Regulation Goal

Figure 4: Interpersonal Emotion Regulation Goals

According to its preconditions, it becomes available for an agent A when A thinks that another agent B likes him less than he likes B. Inspired in Balance Theory, the rationale for this goal is that agent A will try to make the relationship more consistent by increasing in a proportional way B's relation towards him. For this reason, the goal state is for B to have a positive emotion with intensity given by the difference between A's attraction towards B and B's attraction towards A (the variable $[deltaL]$). The rationale here is that the perfect consistency is given by exactly reciprocal relations. If the difference in relations is high, agent A will have to do something really nice to increase B's attraction to a reciprocal level. If the difference is low, agent A could get off with something simple.

In the Reactive Positive Regulation goal, the agent does not take on its own the initiative to change the relation. Instead, the goal is activated in response to the situations when another agent is feeling a negative emotion and the relation towards that agent is positive. The desired goal state is to make the target agent feel a positive emotion, with the same intensity as the negative emotion. Succeeding in this goal will also increase the relation accordingly.

The strategies or actions that can be used to perform interpersonal emotion regulation depend on the target and on the current situation. Any action that is appraised as positive by the target can eventually be used to create a positive emotion in him. Here we follow a similar categorization as used by Niven, and divide actions into two main groups: instrumental and relationship-oriented. Instrumental strategies focus on providing practical and instrumental help to the target, i.e. doing things that are useful to the target. For example, if an agent is upset because of being injured, healing him will be perceived and appraised as good for him and thus create a positive emotion. Instrumental strategies available are pretty much dependent on the domain where the architecture is applied to: the actions defined in the domain, and their effects in terms of appraisal for the agents.

In many situations however, it will not be possible to use instrumental strategies to create emotions in others. As in the case where the target is sharing a past distressful episode, and it is not possible to change that situation. However, it will still be possible to use relationship-oriented strategies, which focus on fulfilling other's socioaffective needs by providing them social support and reassurance. The advantage of these strategies is that they are quite

Table 1: Generic relationship-oriented emotion regulation strategies

Strategy	Preconditions	Description
Validation	the agent really likes the target	making it clear you care about the target
Advice	target is distressed because of a past event	giving advice to the target about the event
Praise	target is happy about a past action he did	praise the target's work
Compliment	none	generic compliment, making the the target feel special.
Social Invitation	the agent really likes the target	inviting the target to get some drinks or to another social event

generic, and can be easily used across several distinct domains. To include these type of strategies in the proposed architecture, a set of generic relationship-oriented actions was defined based on the relationship-oriented strategies pointed out by Niven's study (see table 1). These actions are characterized for having a set of generic preconditions that dictate when they can be used.

When there is more than one alternative strategy possible, a relevant question arises: which strategy should we choose? The idea here is not to generate the strongest emotion in the target, but to select actions that bring the target closer to a desired emotion intensity. To that end, equation 2 is used to determine the distance between the desired appraisal values and the estimated appraisal value for a given action a and a particular target agent t.

$$\Delta Appraisal(V, a, t) = \sum_{V} |v_d - v_{a,t}| \quad (2)$$

V is the set of of appraisal variables used to generate the desired emotion. In the case of OCC, if the aim is to create a Gratitude emotion in a target, we are interested in searching for actions with positive desirability and positive praiseworthiness. $|v_d - v_{a,t}|$ is the distance between the desired value for the appraisal variable v and the estimated appraisal value of variable v for action a according to the target agent t (this estimation is performed by using the Theory of Mind Mechanism to simulate the appraisal of action a from agent t's perspective). The alternative operators generated by the instantiation of the AppraisalDerivation operator are thus ranked according to the function defined. The one with the lowest value, i.e the one with minimum distance between the desired and the expected appraisal, is the one selected to continue planning. However, it does not necessarily mean that it will be the one executed. If the planner is not capable of building a complete valid plan with the selected operator, it will try the next lowest-ranked alternative.

5. CASE STUDY

The model presented was integrated on top of a well-known Role-Playing Game, NeverWinter Nights 2 (NWN2). Our aim was to provide a game-like environment where players have the chance of interacting socially with other characters while pursuing a given task. Moreover, we believed that using a Role-Playing game as an interaction context with the user would facilitate the creation of elaborate social scenarios, in which the development of relations with distinct

characters would take a preponderant role. A scenario was then designed to help us evaluate aspects of the behaviour generated by the model. In particular, we wanted to evaluate the use of interpersonal emotion regulation strategies in the development of a socio-affective relation with the user. To that end, a simple storyline quest was created in NWN2: three seasoned warriors were summoned by the King to kill an evil dragon that is terrorizing the kingdom. The task at hand will require the party to complete several sub quests, fighting numerous foes, to find where the dragon is hiding and find the weapon that can kill him. The user controls the main warrior hero, while the other characters will each be controlled by distinct agents. The scenario was designed so that the user could achieve its quest in 30 minutes, independently of any help received by its teammates.

In this particular scenario we used OCC appraisal theory together with an appraisal mechanism based on drives. Agents have three main drives, Energy, Integrity and Affiliation. Positive social interactions increase affiliation, getting hurt decreases integrity, while getting healed will increase it back again. Energy is spent with combat actions and gained with resting. The appraisal derivation function corresponds to determining an action's desirability by calculating the impact of the action in the agent's drives. This variable is then used to generate emotions according to the OCC theory as described previously.

Two versions of the same scenario were created. In the first one, we used the model without the full emotional intelligence features (i.e the agent is not able to reason about and regulate emotions). Yet, in this version, agents still have emotions and goals such as fighting enemies, and healing themselves when hurt. In the second version, we used the full version of the model including the capability to reason about emotions and perform interpersonal emotion regulation.

In addition to the relationship-oriented strategies depicted in Table 1, an additional instrumental strategy was created. Whenever the player has low integrity, an action can be used to give an healing potion to the user increasing its integrity back. One of the characters (Varsuvius) was designed to have a strong affiliation for the user, meaning that he will try to befriend the player frequently. Figure 5 shows an example of a validation strategy being performed by this character (wearing black) towards the user (wearing shiny armor). The second character has a neutral relation towards the user.

There were however some significant simplifications in the scenario. The most relevant one was that the user was not allowed to take the initiative to interact with the agents, or directly perform any social action that affected them. This is mainly due to the fact that we wanted all users to have (roughly) the same experience. If we would allow the player to affect others, he could change other's behaviour towards him (for instance making the other characters dislike him), thus having an uncontrolled effect in the evaluation process.

6. EVALUATION

Evaluating a computational model as this one is a difficult task. So, it is important to point out that we do not consider to be viable to directly validate the model with user studies in scenarios where the model is implemented and used. Our argument is that applying our model to a particular scenario will often result in a series of simplifications and design choices (that are not part of the model), which will likely have an impact on the perceptions of the users (e.g. selecting which particular utterance will be spoken by a character when performing a relationship-oriented strategy). Nevertheless, performing a user study can be quite useful, because it allow us to understand the effects and the relevance of the type of behaviour generated by the model, and can guide us in further developing our work. With this said, we will now report on the evaluation done measuring the effects of using interpersonal emotion regulation strategies in the development of a friendship relation between Varsuvius and the player.

Figure 5: Validation strategy being applied towards the user

Procedure

Twenty two subjects, 17 male and 5 female, with ages ranging from 20 and 35 years old, took part in the experiment. Most of them were undergraduate or graduated students and were recruited via email. Participants were instructed to play the game and were told they would be questioned about the game experience at the end of the game. Participants were randomly assigned to one of the two versions. At the end of the game, participants were asked to fill a questionnaire about the relationship with Varsuvius, answering from the perspective of the character they were playing. After filling the questionnaire, they were rewarded with a movie ticket and the experiment was over.

Measures

In order to access the friendship level established with the user, we have used McGill Friendship Questionnaire (MFQ) [14], which measures friendship quality. MFQ is composed by two questionnaires, the first one measures the degree to which a friend fullfils the following six friendship functions: (1) stimulating companionship - doing enjoyable or exciting things together; (2) help - providing guidance and other forms of aid; (3) intimacy - being sensitive to the other's needs and states and being open to thoughts and feelings; (4) reliable alliance - remaining available and loyal; (5) self-validation - helping the other maintain a positive self-image; (6) emotional security - providing confidence in novel or threatening situations. The second questionnaire taps respondents' positive feelings for a friend and friendship satisfaction. These two questionnaires were combined into a final questionnaire that contains a set of items for each one of the six friendship functions, satisfaction and positive feel-

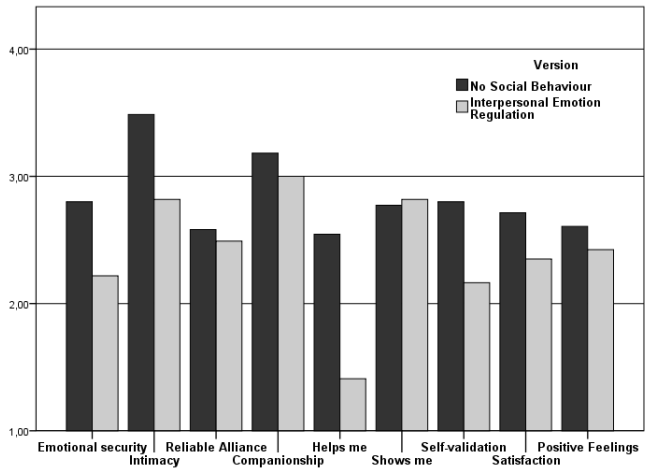

Figure 6: Results obtained for the friendship functions, satisfaction and positive feelings (1 - Totally Agree, 5 - Totally Disagree).

ings. Participants express their agreement or disagreement about each assertion using a five-point Likert scale (1 - Totally Agree; 5 - Totally Disagree).

Results

We started by performing a Cronbach alpha test to evaluate the internal consistency of the six friendship functions, satisfaction, and positive feelings. All the dependent variables had at least satisfactory consistency except for the help friendship function. Therefore we divided the help items into two groups with good consistency: one for questions such as "Varsuvius helps me when needed" (helps-me), and another for questions such as "Varsuvius shows me how to do things better" (shows-me). Next, we applied a Mann-Whitney U test and found statistically significant effects for the following variables: self-validation - $U=26$, $p<0.022$, $r=-0.49$(medium); helps-me - $U=9.5$, $p<0.01$, $r=-0.74$(strong); emotional security - $U=24$, $p<0.016$, $r=-0.52$(strong); intimacy - $U=23.5$, $p<0.015$, $r=-0.52$(strong).

Figure 6 shows the averages for the dependent variables measured (the smaller the value, the better). Four of the six friendship functions yielded statistically significant results, indicating that the emotionally intelligent agent who performed interpersonal emotion regulation was perceived as providing more intimacy, help, emotional security and self-validation functions.

Discussion

The results obtained for intimacy, help-me, emotional security and self-validation suggest that, the use of interpersonal emotion regulation strategies makes agents to be perceived as more friendly. We also believe that there might be different effects by using instrumental or relationship-oriented strategies. For example, the helps-me variable is likely related to the use of instrumental strategies (in this case giving potions), while self-validation is likely related to relationship-oriented strategies. Our initial idea was to determine the overall influence of Interpersonal Emotion Regulation in friendship functions, but in the future we could test these two types of strategies independently.

Although the average values obtained for the remaining variables (companionship, reliable alliance, shows-me, satisfaction and positive feelings) are slightly better in the version with Emotional Intelligence, there were no statistically significant differences between the two versions. There are several issues that can explain a lower or non-existing effect on satisfaction and positive feelings. First and foremost, the measure used was designed to address the relationship between two close friends. In our scenario, the agents do not create a relation with the user but with the character that the user is playing in the game. This could explain why users do perceive the agent as being more friendly, but do not have a similar strong increase in friendship satisfaction. Secondly, the fact that the user is not really able to respond emotionally back to the agent can also have an impact in these results. It may be that creating a reciprocal strong relationship between user and agent requires a bidirectional interaction in terms of social behaviour.

7. CONCLUSION AND FUTURE WORK

This paper addressed the problem of creating agents are able to intentionally establish and strengthen social relations with others. Inspired in studies that show that emotionally intelligent individuals are more liked by others, our approach focused on modeling emotional intelligence capabilities, by endowing the agents with the possibility of reasoning about emotions of others, and applying interpersonal emotion regulation strategies, i.e. creating emotions in others. We proposed a generic model for emotionally intelligent agents that is able to generate emotions, reasons backwards about emotions by using two emotional meta-operators, and performs interpersonal emotion regulation. It also has an an explicit model of social relations that evolves according to the experienced emotions. The model works with a set of of domain-independent emotion regulation goals to help the agent decide when to perform Interpersonal Emotion Regulation. Further, the model has a set of generic relationship-oriented regulation strategies that can be used across distinct domains.

Using this innovative model, a scenario was created on top of NWN2. One evaluation was performed to test if the social behaviour generated by the agents could indeed lead to the establishment of stronger affective relations with the user in an interactive narrative or game environment. By applying McGill Friendship Questionnaire (MFQ) in a study with twenty two participants we were able to show that when one of the agents employed interpersonal emotion regulation strategies, resulting from the architecture developed, that agent was perceived as more friendly, having a better score on 4 out of 6 friendship functions (intimacy, self-validation, help and emotional security).

Unfortunately, we were not able to detect an effect on friendship satisfaction nor positive feelings. This means that the player's relation towards the agent did not seem to increase much in terms of satisfaction. We suspect that the limited interactivity from the player towards the agents may have caused this effect. Natural social interactions are often very reciprocal, if someone does something nice to you, you are expected to retribute. This hypothesis will require further tests in future work.

Although the results achieved with the evaluation are promising, one must acknowledge that the model proposed was not directly validated with the performed user study. What

was tested was the fact that running this architecture, in real time, in our characters, allows the generation of character's behaviours that lead to the establishment of different social relations with the user. Nonetheless, this does not remove the importance from the user study or from the proposed model. In fact, this work allowed us to show that employing interpersonal emotion regulation strategies can indeed help us establish some aspects of friendship relations with synthetic characters in different ways. And this is of particular importance for the virtual agents community, because it will lead to an increase flexibility in the creation of these types of agents. We can easily build intelligent virtual agents that, by simple authoring change the relation they establish with the users.

Finally, one argument that is often raised against the creation of agent architectures/agent models for characters in games or virtual environments is that we can achieve the same, or even more "believable" behaviour if we script all the actions of the characters in those environments. Often people argue that scripted characters are easier to build, specially if we want to guarantee certain degree of believability. Yet, scripted characters do not have the flexibility to change, adapt and re-use that we can find in autonomous agents models for virtual environments. In a similar way as procedural content generation is being built for games, we believe that the creation of complex autonomous behaviours for agents and NPCs will foster the development in the area of intelligent virtual agents and allow for stronger and more socially rich interactions with these characters in the long run. Furthermore, the model can easily be applied to a different context/scenario, without needing to respecify any of the Interpersonal Emotion Regulation Goals and the relationship-oriented strategies. The only requirement is to specify as planning operators the available instrumental actions in the scenario, and to map some of the social actions to particular utterances. By doing this, emotionally intelligent behavior will automatically emerge from the model.

As future work, we plan to create slightly different scenarios to try to follow up on the problems identified by this initial evaluation. Another aspect worth investigating is the role of instrumental strategies vs relationship-oriented strategies, which can be done by isolating them in each scenario.

Acknowledgments

This work was partially supported by the European Community (EC), through the EU FP7 ICT-215554 project LIREC, and by national funds through Fundação para a Ciência e a Tecnologia (FCT), under project PEst-OE/EEI/LA0021/2011 and under a scholarship (SFRH BD/19481/2004). 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] C. Adam. "Emotions: from psychological theories to logical formalization and implementation in a BDI agent". PhD Thesis. Institut National Polytechnique de Toulouse, 2007.

[2] R. Aylett, J. Dias, and A. Paiva. "An affectively driven planner for synthetic characters". In: *Proceedings of International Conference on Automated Planning and Scheduling ICAPS06*. UK, 2006.

[3] R. Aylett et al. "But that was in another country: agents and intercultural empathy". In: *Proceedings of Autonomous Agents and Multiagent Systems (AAMAS)*. ACM Press, 2009.

[4] T. Bosse and F. Lange. "Development of Virtual Agents with a Theory of Emotion Regulation". In: *Proceedings of the Eight IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT08*. IEEE Computer Society Press, 2008, pp. 461–468.

[5] J. Dias et al. "Lie to Me: Virtual Agents that Lie". In: *to appear in Proceedings of Autonomous Agents and Multiagent Systems (AAMAS)*. Minnesota, USA, 2013.

[6] S. Epstein. "The self-concept, the traumatic neurosis, and the structure of personality". In: *Perspectives on personality* 3A (1991), pp. 63–98.

[7] R. E. Fikes and N. J. Nilsson. "STRIPS: A new approach to the application of theorem proving to problem solving." In: *Artificial Intelligence* 2.3-4 (1971), pp. 189–208.

[8] F. Heider. *The Psychology of Interpersonal Relations*. NY: Wiley, 1954.

[9] Mayer J., Caruso D., and P. Salovey. "Emotional intelligence meets traditional standards for an intelligence". In: *Intelligence* 27.4 (1999), pp. 267–298.

[10] P. Lopes et al. "Emotion Regulation Abilities and the Quality of Social Interaction". In: *Emotion* 5.1 (2005), pp. 113–118.

[11] P. N. Lopes, P. Salovey, and R. Straus. "Emotional intelligence, personality and the perceived quality of social relationships". In: *Personality and Individual Differences* 35 (2003), pp. 641–658.

[12] S. Marsella, J. Gratch, and P. Petta. "A blueprint for an affectively competent agent: Cross-fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing." In: Oxford University Press, 2010. Chap. Computational Models of Emotion.

[13] S. Marsella, L. Johnson, and C. LaBore. "Interactive Pedagogical Drama". In: *In Proc. of the Fourth International Conference on Autonomous Agents*. ACM Press, 2000, pp. 301–308.

[14] M. Mendelson and F. Aboud. "Measuring friendship quality in late adolescents and young adults: Mcgill friendship questionnaires". In: *Canadian Journal of Behavioural Science* 31.1 (1999), pp. 130–132.

[15] J. Moreno. *Who shall survive? : a new approach to the problem of human interrelations*. Washington D.C.: Nervous and Mental Disease Publishing Co., 1934.

[16] B. Mott and J. Lester. "Narrative-Centered Tutorial Planning for Inquiry-Based Learning Environments". In: *Proceedings of Intelligent Tutoring Systems*. Springer, 2006.

[17] K. Niven, P. Totterdell, and D. Holman. "A Classification of Controlled Interpersonal Affect Regulation Strategies". In: *Emotion* 9.4 (2009), pp. 498–509.

[18] M. Ochs, N. Sabouret, and V. Corruble. "Simulation of the Dynamics of Non-Player Characters' Emotions and Social Relations in Games". In: *IEEE Transactions on Computational Intelligence and AI in Games* 1.4 (2009), 281Ű297.

[19] A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. UK: Cambridge University Press, 1998.

[20] D. Pizzi and M. Cavazza. "Affective Storytelling based on Characters' Feelings". In: *Proceedings of the AAAI 2007 Fall Symposium on Intelligent Narrative Technologies*. Arlington, Virginia, 2007.

[21] R. Prada and A. Paiva. "Teaming Up Human with Synthetic Characters". In: *Artificial Intelligence* 173.1 (2009), pp. 257–267.

[22] D. Pynadath. "PsychSim: Agent-based Modeling of Social Interactions and Influence". In: *ICCM*. 2004, pp. 243–248.

[23] P. Salovey and J. Mayer. "Imagination, Cognition and Personality". In: 1990. Chap. Emotional Intelligence, pp. 185–211.

[24] S. Taylor and J. Brown. "Illusion and well-being: A social psychological perspective on mental health". In: *Psychological Bulletin* 103 (1988), pp. 193–210.