Development of Virtual Agents with a Theory of Emotion Regulation

Tibor Bosse and Frank P. J. de Lange
Vrije Universiteit Amsterdam, De Boelelaan 1081a,
1081 HV Amsterdam, The Netherlands
{tbosse, fle600}@few.vu.nl

Abstract

In order to endow virtual agents with more realistic affective behavior, it is important to provide them not only the capability to generate and regulate emotions, but also the ability to reason about the emotion regulation processes of other agents. To this end, this paper introduces a computational model for a Theory of Emotion Regulation (ToER). The model has been implemented and tested using the modeling language LEADSTO. In addition, a virtual environment application has been developed, which is inhabited by agents that are equipped with the model for ToER. A first evaluation indicates that the model indeed enables the agents to show more realistic affective behavior.

1. Introduction

Over the past years, there has been an increasing interest in the application of *intelligent virtual agents* in various domains [19]. Typical examples are conversational agents (e.g., representatives on the internet), agents in (serious) games (e.g., an instructor in a flight simulator), and agents in computer-generated virtual stories. Recently, much research has been dedicated to developing virtual agents with more realistic graphical representations. However, the affective properties of such agents are usually rather limited, and not very human-like. For example, although many IVAs currently have the ability to somehow show emotions by means of different facial expressions, it is quite difficult for them to show the right emotion at the right moment. One step further, it is even more difficult for them to actually *understand* and react empathically to the emotional state of other agents. This is in conflict with the requirement of virtual agents to closely mimic human affective behavior. Several studies in Social Sciences have shown that this is an important prerequisite for an agent to increase human involvement in the virtual environment [15]. Therefore, existing systems based on IVAs are not as effective as they could be. Some properties that they typically lack are the ability to show emotions (not only in terms of facial expression, but also in terms of behavior), and to have insight in each other's mental states.

Within the last decade, the literature in Agent Systems shows an increasing amount of attempts to incorporate emotions within software agents (e.g., [22]). In general, two classes of approaches can be distinguished: those that are trying to equip agents with *emotion elicitation* processes (e.g., [3]) and those focusing on agents with *emotion regulation* (or coping) processes (e.g., [6], [17]). The first process addresses the way how agents develop emotions, based on stimuli from the environment (e.g., [12]), whereas the second process addresses the way how agents control their emotions in case they do not correspond with the emotions they desire to have (e.g., [14]).

When IVAs start to show more realistic affective behavior, the interaction between such agents may also become more realistic. More specifically, agents may start to get insight in the emotional state of the agents around them, and try to influence those. For example, a player in a virtual soccer game that notices that his teammate is losing confidence may try to encourage him. To incorporate these kinds of capacities into IVAs, they will need the capability to attribute mental states to others, and to reason about these. In psychology, this characteristic is often referred to as *Theory of Mind* (or ToM, see, e.g., [2]). According to [5], agents can exploit a ToM for two purposes: to *anticipate* the behavior of other agents (e.g., preparing for certain actions that the other will perform), and to *manipulate* it (e.g., trying to prevent the other from performing certain actions).

A number of approaches in the literature address the development of virtual agents with a ToM, e.g., [5], [18]. Usually, such agents maintain, in one way or the other, a model of the epistemic (e.g., beliefs) and/or

978-0-7695-3496-1/08 $25.00 © 2008 IEEE
DOI 10.1109/WIIAT.2008.65

motivational states (e.g., desires, intentions) of other agents. However, since the latest generation of IVAs is starting to have emotions, such agents ideally also have insight in each other's emotions. This idea is in line with the theories of many Cognitive Scientists like Gärdenfors, who claims that humans have a ToM that is not only about beliefs, desires, and intentions, but also about other mental states like emotional and attentional states [13]. Based on these ideas, this paper addresses the following challenge:

Is it possible to develop IVAs with a Theory of Emotion Regulation (ToER), i.e., a ToM that also considers emotional states, and more specifically, emotion regulation, which allows them to reason about each other's emotion regulation processes?

In Section 2, first a generic model for such a ToER will be described at a conceptual level, using the modeling language LEADSTO [4]. In Section 3, a number of simulation results are shown that were generated based on this model, in order to test its basic mechanism. In Section 4, we will introduce a virtual environment application that is inhabited by virtual agents with the model for ToER. Finally, in Section 5 the model is evaluated and conclusions are drawn.

2. Conceptual Model

In this section, the model for ToER will be described at an intuitive, conceptual level, using the agent-based modeling language LEADSTO [4]. This language allows the modeler to integrate both qualitative, logical aspects and quantitative, numerical aspects. In LEADSTO, direct temporal dependencies between two state properties in successive states are modeled by executable dynamic properties. The format is defined as follows: let α and β be state properties of the form 'conjunction of ground atoms or negations of ground atoms'. In LEADSTO the notation $\alpha \twoheadrightarrow_{e,\, f,\, g,\, h} \beta$ means:

If state property α holds for a certain time interval with duration g, then after some delay (between e and f) state property β will hold for a certain time interval of length h.

Here atomic state properties can have a qualitative, logical format, such as an expression desire(d), expressing that desire d occurs, or a quantitative, numerical format such as an expression has_value(x, v) which expresses that variable x has value v. For more details, see [4].

Below, the LEADSTO model for ToER is introduced in a number of steps[1]. First, some basic mechanisms to model an agent's reasoning in terms of beliefs, desires and intentions (BDI), and its emotion regulation processes are introduced. Next, these ideas are combined to obtain a model for emotions and BDI (EBDI). After that, a basic mechanism to model an agent's Theory of Mind (ToM) is introduced, and finally this model is combined with the EBDI model, to obtain a model for ToER.

2.1. BDI-modeling

The BDI-model bases the preparation and performing of actions on beliefs, desires and intentions (e.g., [21]). This model shows a long tradition in the literature, going back to Aristotle's analysis of how humans (and animals) can come to actions; cf. [1]. He discusses how the occurrence of certain internal (mental) state properties within the living being (such as desires) causes the occurrence of an action in the external world.

More recently, researchers in the domain of Agent Technology have used these ideas as a basis to develop the well-known BDI-model of reasoning [21]. This model forms the foundation of the model for ToER presented in this paper. See Figure 1 (for the time being, only focus on the part below the dotted line) for the generic structure in causal-graph-like style, as often used to visualize LEADSTO specifications. Here the box indicates the borders of the agent, the circles denote state properties, and the arrows indicate dynamic properties expressing that one state property leads to (or causes) another state property. In this model, an *action* A is performed when the agent has the *intention* I to do this action and it has the *belief* that certain circumstances in the world are fulfilled such that there is an opportunity to perform the action. Beliefs are created on the basis of observations. The intention to do a specific type of action is created if there is some *desire* D, and the agent believes that performing this action will fulfill this desire. Such relations within the general BDI-model can be specified in formal LEADSTO format as follows (where the timing parameters e, f, g, h have been omitted for simplicity):

desire(D) $\wedge$ belief(satisties(A, D)) $\twoheadrightarrow$ intention(A)
intention(A) $\wedge$ belief(opportunity_for(A)) $\twoheadrightarrow$ performs(A)

with appropriate desire D and action A. Note that the beliefs used here both depend on observations, as

[1] See file IVA08-simulation.lt at [23] for the complete LEADSTO specification.

shown in Figure 1. Furthermore, $\wedge$ stands for the conjunction operator (and) between the atomic state properties (in the graphical format denoted by an arc connecting two (or more) arrows).

2.2. Emotion Regulation

Gross [14] describes the process of emotion regulation as “all of the conscious and nonconscious strategies we use to increase, maintain, or decrease one or more components of an emotional response”. A simple example of an emotion regulation strategy is closing your eyes when you are watching an exciting penalty shootout. To formally represent this principle within our IVAs, we reuse parts of the difference equation model by [7], as explained below.

A first design decision, inspired by [7], is to concentrate on one specific type of emotion at a time. In principle, this can be any emotion that is considered to be a basic human emotion, e.g., anger, happiness, or fear [10]. The *emotional response level* (ERL, see the upper part of Figure 1) that an agent experiences for such an emotion is modeled by a real number in the interval [0, 2] (where 0 is the lowest possible ERL, and 2 the highest). Next, this ERL may be influenced by the observation of certain *events* in the environment (i.e., emotion elicitation). Each event has a number of attributes (number between 0 and 1) that determine to what extent they influence a certain emotion. For example, an encounter with a fierce animal may increase the ERL for fear with 0.8. In addition, for each agent a *baseline ERL* is assumed[2]. This is a specific ERL that an agent (either consciously or unconsciously) tries to achieve, for a certain emotion. For instance, some people who enjoy extreme sports will have a high baseline ERL for excitement, while others, who prefer a more quiet kind of recreation, will have a lower baseline ERL.

Finally, the process of emotion regulation is modeled by introducing a parameter β for *regulation speed*. This is an individual parameter for each agent, which determines how quickly the agent regulates its ERL towards its baseline ERL. For example, after an unpleasant experience, some agents will stay angry for a long time, while others will leave the event behind very quickly [14]. The regulation process is modeled via the following LEADSTO rule:

ERL(X) $\wedge$ baseline_ERL(Y) $\wedge$ reg_speed(β) $\twoheadrightarrow$
ERL(X+(Y-X)*β)

[2] This concept is comparable to [7]’s notion of “Emotional Response Level Norm”.

2.3. EBDI-modeling

A next step is combine the model for emotion regulation with the BDI-model presented earlier. In the cognitive literature, it is often claimed that cognition can be divided into two distinct systems: a low-level, emotional and unconscious system, and a high-level, evolutionary recent, conscious system, see, e.g., [11]. According to this view, an important role of emotions is to function as some kind of learned shortcut, allowing a person to quicker come to certain actions (e.g., [9]). Such theories have formed the basis of recent attempts to develop computational models that integrate emotional and rational concepts, e.g., [16], [20].

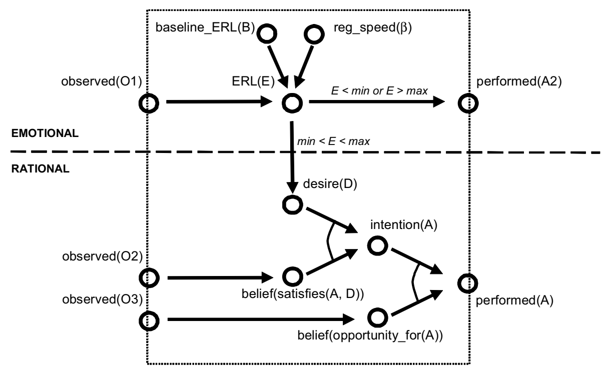

Fig. 1. Generic structure of the EBDI-model

In line with these ideas, we propose to integrate the two models introduced in the previous subsections as follows (see Figure 1). An agent usually reasons rationally, according to the BDI-model (lower part). Meanwhile, at any time point, it has a certain ERL, which is influenced by observations of events, and by its own emotion regulation processes, based on its baseline ERL (upper part). As long as the ERL stays within certain boundaries min and max, it only has a marginal impact on the agent’s actions. It does influence its reasoning process (by influencing the desires the agent has, see the downward arrow), but the agent keeps on reasoning rationally. However, whenever the ERL becomes lower than min or higher than max, the reasoning process is bypassed, and the agent acts emotionally. For example, an agent that is walking through a dark forest usually stays close to its friend, but whenever it becomes too frightened it might start panicking.

2.4. Theory of Emotion Regulation (ToER)

So far, a model was presented that describes the cognitive processes of a single agent. To make this model the subject of another agent's theory of mind, the idea of *recursive modeling* is used [18]. This means that the beliefs that agents have about each other are represented in a nested manner. To this end, each mental state is parameterized with the name of the agent that is considered, creating concepts like desire(agentA, D) and ERL(agentB, 0.75). In addition, a number of meta-representations, expressed by meta-predicates are introduced. For example, belief(agentA, ERL(agentB, 0.5)) states that agent A believes that agent B has an ERL of 0.5. Similarly, desire(agentA, not(performed(agentB, A1))) states that A desires that B does not perform action A1.

In addition to such meta-statements about *state properties*, meta-statements about *dynamic properties* are used, using the depends_on relation. For example, the statement belief(agentA, depends_on(performed(agentB, A1), intention(agentB, A1))) indicates that agent A believes that, whether agent B performs action A1 depends on whether B intends to do A1. In such a way, an agent may have knowledge about a complete EBDI-model (as depicted in Figure 1) of another agent.

Based on such knowledge, an agent may reason through the EBDI-model of another agent in various manners. First, types of simple forward reasoning may be used, e.g., by using the following LEADSTO rule:

belief(A, X) ∧ belief(A, depends_on(Y,X)) ↠ belief(A, Y)

For example, if A believes that B has a medium emotional state, and A observes that B experiences a very emotional event, then A will believe that B´s emotional state will increase. In addition, several backward reasoning strategies may be used, e.g.:

desire(A, Y) ∧ belief(A, depends_on(Y,X)) ↠ desire(A, X)

For example, if A desires that B does not run away, but believes that B will run away if it becomes too frightened, then A will desire that B does not become too frightened.

2.5. Updating the ToER

In the domain of virtual environments, obviously, the model an agent has about another agent is not always completely correct. Therefore, an agent should keep on updating its model about the other whenever possible. For the case of the ToER, the main concepts involved in the model are the *ERL*, the *baseline ERL*, and the *regulation speed*. Our IVAs estimate the values of these concepts in a way that is similar to the way humans estimate each other's personality aspects. The only information that an IVA has to update these values is provided by the externally observable behavior of the other agent. For example, if agent A believes that agent B has a high ERL (e.g., is very frightened), and believes that this will lead to a certain action (e.g., screaming), but observes that B does not perform this action, then A will lower its estimated ERL of B. A more generic variant of his process is represented by the following LEADSTO rule (where δ is the extent to which the estimated ERL is decreased):

belief(A, ERL(B, X)) ∧ observed(A, not(performed(B, A1)) ∧
belief(A, depends_on(performed(B, A1), ERL(B, X)) ↠
belief(A, ERL(B, X-δ))

In addition to updating the ERL, our agents use two other strategies to update the ToER. Because of space limitations, the formal details of these strategies are not shown here. However, they are summarized below in pseudo-code:

Strategy 1 - Update Estimated Baseline
"If the agent shows 'positive'[3] emotions more often than expected, and 'negative' emotions less often than expected, then increase the estimated baseline ERL." (and vice versa)

To illustrate the mechanism of this strategy, consider Figure 2, which shows examples of an actual and an estimated emotion curve. In both graphs, time is on the horizontal axis, and the ERL is on the vertical axis. The solid lines indicate the dynamics of the ERL. These dynamics are influenced by the occurrence of events (which are the same for both graphs): first an emotion increasing event occurs, and later an emotion decreasing event (indicated by the arrows). Moreover, in both graphs emotion regulation takes place after an event has occurred, but the value of the regulation speed β is lower in the left graph (the curves are less steep). Furthermore, the baseline ERL towards which the emotion is regulated (indicated by the dotted line) is higher in the left graph.

Clearly, Figure 2 illustrates a situation in which the estimated model does not match the reality very well. Thus, the ToER has to be updated. To this end, as mentioned earlier, the agent with the ToER will make use of the externally observable behavior of the other agent. To make this possible, it is assumed that each agent has a certain range of ERLs in which its behavior is not much driven by emotions. This range is indicated by the shaded area in Figure 2; its borders correspond to the min and max values in Figure 1. If the ERL ends

[3] The words 'positive' and 'negative' are used in the mathematical sense here. Thus, a positive emotion is an emotion that corresponds to a high ERL (e.g. a high level of fear), and vice versa for a negative emotion.

up outside of this area, the agent in question will show emotional behavior, such as screaming or (opposite case) acting over-confidently.

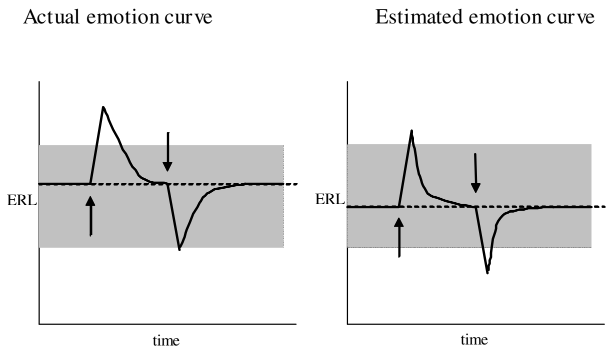

Fig. 2. Actual and Estimated Emotion Curves

These emotional expressions (indicated by the parts of the graphs of Figure 2 outside the shaded area) are used by the estimating agent to update its ToER. For example, Strategy 1 states that, if the graph of the actual situation is too often above the shaded area compared to the estimated situation (which is the case in Figure 2, e.g., the agent panics more often than expected), then the estimated baseline ERL should become higher. In addition to Strategy 1, another important strategy is:

Strategy 2 - Update Estimated Regulation Speed
"If the agent shows each of its emotions for a longer period than expected, then decrease the estimated regulation speed β."

This strategy is meant to fine-tune the ToER. Therefore, it only becomes useful after a while, when a more or less solid estimation of the baseline ERL has been made[4]. The strategy can be explained by referring to Figure 2 as well. Clearly, if the actual graph lies outside the shaded area for much longer periods than the estimated graph, the value of β should be decreased.

2.6. Manipulation based on ToER

As soon as an agent has established a reasonable model of another agent, it may use this to make predictions about this agent's behavior, using forward reasoning strategies as introduced earlier. In addition, it may reason about which behavior of the other agent it does not desire, using backward reasoning strategies. Based on these two types of reasoning methods, an agent may deliberate about how to *manipulate* another agent in such a way that that agent satisfies its desires. For example, suppose an agent knows that its friend is very scared, and it also knows that some scaring event is about to happen, which will make the other so scared that he will run away, which it does not want. Then, the agent may decide to calm down its friend, e.g., by talking to him. This principle is modeled via the following LEADSTO rule:

desire(A, ERL(B, X)) ∧ belief(A, ERL(B,Y)) ∧ X<Y ↠
desire(A, decrease_ERL_by(B,ε))

Here, the parameter ε indicates the extent to which the ERL of the other should be decreased. This may be either a constant value or a function of X and Y. Similar rules have been created to determine when to increase the ERL.

3. Simulation Results

To test the basic mechanism of the model for ToER, it has been used to perform a number of simulation runs within the LEADSTO simulation environment [4]. The model was tested in a small scenario, involving two agents, A and B. Agent A was equipped with the model for ToER, whereas agent B was not. The central emotion used in the scenario is fear. In order to simulate this, every now and then certain events take place, which influence the level of fear of the agents either positively (e.g., an encounter with a dangerous creature) of negatively (e.g., an encounter with a friendly person). The main goal of agent A is to continuously estimate the level of fear of the agent B, and to manipulate this agent when necessary. To this end, agent A starts with some default model of agent B, and then keeps on updating this using the strategies explained earlier. When agent B becomes too frightened, agent A has to calm it down; similarly it has to ensure it does not become too over-confident (e.g., by sometimes reminding it that dangerous events may still happen).

Figure 3 shows an example of (part of) a resulting simulation trace. Time is on the horizontal axis, and different state properties are on the vertical axis. A box on top of a line indicates that a state property is true. The upper three lines represent the observations (enemies and friends) and actions (screaming) of agent B. The next two lines indicate agent B's actual baseline ERL (of 1.2) and regulation speed (of 0.1). After that, two lines about agent A's estimations of B's baseline ERL are shown, and two lines about A's estimations of B's regulation speed. Finally, two graphs are depicted

[4] This is similar to the approach that humans take when they estimate someone's character: first they make a rough estimation (e.g., "how often does he get angry?"), later they start to learn more details (e.g., "how long does it take before he calms down again?").

that represent respectively, B’s actual ERL, and A’s estimations about B’s ERL.

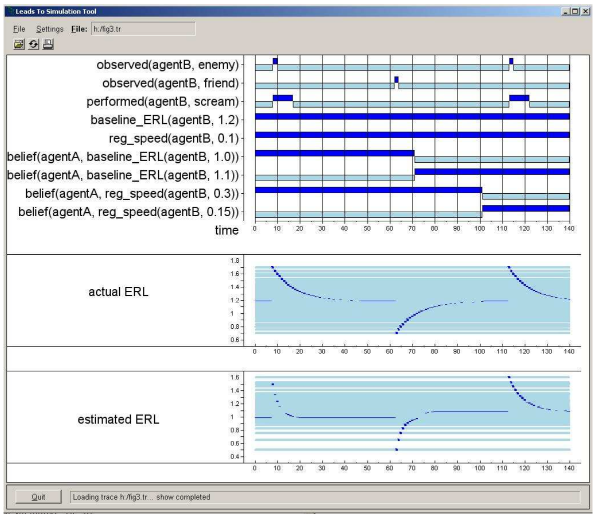

Fig. 3. Example LEADSTO Simulation Trace

As can be observed, over time agent A’s estimations about agent B become more accurate. Agent A uses the strategies described in the previous section to first make a correct estimation about B’s baseline ERL (which clearly should be increased, since B is scared very easily), and later about its regulation speed. As a result, already after three events agent A’s estimations come pretty close to the actual values: 1.1 for the baseline ERL, and 0.15 for the regulation speed β.

All in all, this scenario has been simulated various times, using different starting parameters for ERL, baseline ERL, and regulation speed. Due to space limitations, the details of these simulations are not shown here[5]. The results indicate that the estimation of the baseline ERL never deviates more than 7% from the actual value. The estimation of the regulation speed sometimes has a bigger error margin (varying between 0% and 60%), but this can easily be improved by choosing smaller step sizes (e.g., a smaller δ) in the update formulas of the previous section. In any case, the global patterns shown by the simulated traces are comparable to the way humans estimate each other’s emotions, as described in literature like [13], [14], which is sufficient for the current purposes.

4. Virtual Environment Application

After these initial tests pointed out that the model showed acceptable behavior, it has been implemented within a dynamic virtual environment application[6]. To this end, a graphical environment has been developed (implemented in C++) that is comparable to classical 2D adventure games like “The Legend of Zelda”[7]. Similar to the LEADSTO simulations, this environment is inhabited by two agents, called the *traveler* and the *guide*. The traveler is an IVA that needs to travel through an imaginary world towards a safe destination. This world is inhabited by a various dangerous creatures like dragons and wolves, but there are also friendly people who travel through the area. Moreover, the world is subject to natural phenomena like lightning and earthquakes. These encounters and events occur at random time points, but in more or less fixed areas of the world. Since the traveler has no knowledge of the environment he is in, he has to follows the path of another IVA, the guide. The guide has decent knowledge about the environment. Based on this, the guide helps the traveler to navigate through the world by showing him the way.

However, a major complication for the guide is that the traveler’s behavior may be influenced by emotions (in particular, the emotion of fear). For example, if his level of fear becomes too high, the traveler will run away, and if it becomes too low, the traveler will act recklessly. For these reasons, the guide needs to maintain a correct ToER of the traveler, and to manipulate him if necessary. Fortunately, since the guide is familiar with the environment, he is able to predict roughly whenever certain events will occur. Based on these predictions and the ToER of the traveler, the guide will make sure the traveler reaches the final destination.

The scenario consists of two main stages. In the first stage, the two IVAs travel through the virtual world, and will experience several events, varying from scaring (fierce creatures, natural phenomena) to relieving (other travelers). In this first stage, the guide will attempt to establish an adequate ToER of the traveler, by observing his reactions to these events. The emotional reactions shown by the traveler are either ‘positive’ emotions (screaming, panicking, or running away), or ‘negative’ emotions (sighing, or recklessly walking around), dependent on his ERL. In the second stage, the IVAs still travel through the virtual world, but now the guide’s ToER is sufficiently reliable to be used for manipulation.

A screenshot of the application is depicted in Figure 4. The left-hand side of the screen displays the virtual environment, with the two IVAs and several other

---

5 More details about the simulation results can be found in IVA08-simresults.pdf at [23].

6 The application (and two movies showing an example run) can be downloaded from [23].

7 See http://www.zelda.com/universe/game/zelda/.

virtual characters and objects. Expressions of emotion and manipulative actions by the IVAs are depicted by text balloons. The right-hand side displays two curves: one for the actual ERL of the traveler, and one for the ERL as estimated by the guide. Furthermore, some additional information is shown, such as the actual and estimated values for baseline ERL and regulation speed. Finally, in those situations where the guide deliberates about how to manipulate the traveler (only in the second stage), another window is shown, which displays a graphical representation of the guide's reasoning process (not shown here due to space limitations).

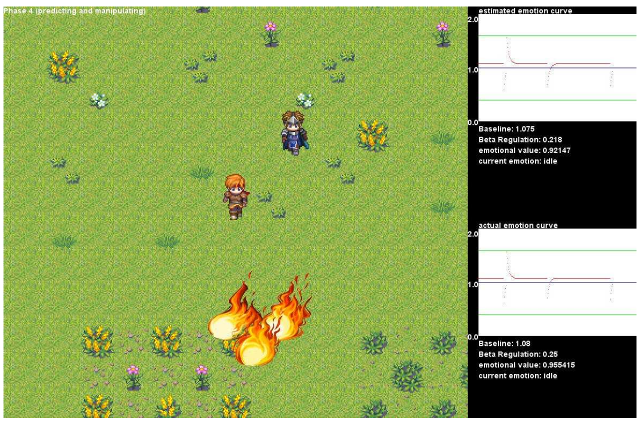

Fig. 4. Screenshot of the Virtual Environment

The application has been tested extensively, using a large amount of different parameter settings (which can be set by hand at the start of the scenario). In all tests, the guide was capable of performing successful manipulative actions at appropriate times. To further evaluate the application, 20 participants were asked to judge the behavior of the IVAs, by answering questions like "do you feel attached to the guide?" and "do you think the capability to estimate emotions makes the guide more realistic?". The initial results of this evaluation were very promising: the participants very much appreciated the guide's capability to estimate and manipulate the affective behavior of the traveler. Nevertheless, in a later stage, a more elaborated evaluation experiment will be performed, in cooperation with colleagues from Social Sciences.

5. Conclusion

To develop virtual agents with more realistic affective behavior, they require not only the capability to generate and regulate emotions, but also the ability to reason about the emotion regulation processes of other agents. To this end, this paper introduces a computational model for a Theory of Emotion Regulation. The model integrates three main components, namely a BDI-model for rational reasoning (cf. [21]), a model for emotion regulation (cf. [6]) that was inspired by Gross's psychological theory [14], and a model for ToM (cf. [18]. Although several approaches in the literature address these topics separately, or integrate two of them, we are not aware of any earlier attempts to synergize all three components. For example, [16], [20], [22] present elaborated models to integrate emotions into the BDI framework, but they do not address the dynamics of emotion regulation in detail, nor do they incorporate these processes into the reasoning of another agent. Other approaches, e.g., [5], [18] do address this topic of reasoning about other agents, but usually have an emphasis on reasoning about epistemic and motivational states, whereas our approach enables an agent to reason about emotional states, and to learn from this.

The model for ToER has been implemented and tested using the modeling language LEADSTO. In addition, a virtual environment application has been developed, which is inhabited by agents that are equipped with the model for ToER. An initial evaluation indicates that the model indeed enables IVAs to show more realistic affective behavior.

The application currently does not contain any interaction with users. Rather, it has the characteristic of a (non-interactive) *virtual storytelling* application, cf. [8]. However, since the conceptual model has been specified in an intuitive, implementation-independent executable format, it is relatively easy to plug it in into any kind of application (e.g., adventure games, training environments, or storytelling applications), including interactive ones.

For future work, it is planned to extend the model in various ways. For example, in addition to emotion regulation, the process of *emotion elicitation* (which is currently modeled by simply adding a fixed number to the ERL, for each observed event) can be modeled in more detail. When doing that, it will be interesting to investigate existing approaches that take *context* into account, e.g. by enabling IVAs to extract certain emotion-eliciting features from the environment. This way, it is also easier to distinguish individual differences in the emotional impact of events. For instance, some agents are particularly afraid of wolves, whereas others might be more afraid of earthquakes. Another possibility for further research is to apply the model for ToER on agents of which the behavioral specification is unknown. Although the mechanisms we used to describe emotional dynamics are rather standard[8], it remains an open question how well the presented strategies work to estimate the behavior of an unknown agent that is not using the expected model.

Finally, an exciting challenge is to apply our model in the domains of Human-Computer Interaction and Ambient Intelligence. Then, the idea would be to enable an intelligent ambient agent to estimate the emotional state of a human user, so that it will be able to support that human in its task. This direction will be further explored in the near future.

6. References

[1] Aristotle (350 BC). *De Motu Animalium.* On the Motion of Animals (translated by A. S. L. Farquharson).

[2] Baron-Cohen, S. (1995). *Mindblindness*. MIT Press.

[3] Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, vol. 37, issue 7, pp. 122-125.

[4] Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2007). A Language and Environment for Analysis of Dynamics by SimulaTiOn. *International Journal of AI Tools*, volume 16, issue 3, pp. 435-464.

[5] Bosse, T., Memon, Z.A., and Treur, J. (2007). A Two-Level BDI-Agent Model for Theory of Mind and its Use in Social Manipulation. In: *Proc. of the AISB 2007 Workshop on Mindful Environments*, pp 335-342.

[6] Bosse, T., Pontier, M., Siddiqui, G.F., and Treur, J. (2007). Incorporating Emotion Regulation into Virtual Stories. In: Pelachaud, C. et al. (eds.), *Proceedings of the Seventh International Conference on Intelligent Virtual Agents, IVA'07*. Springer LNAI, vol. 4722. pp. 339-347.

[7] Bosse, T., Pontier, M., and Treur, J. (2007). A Dynamical System Modelling Approach to Gross' Model of Emotion Regulation. In: Lewis, R.L., Polk, T.A., and Laird, J.E. (eds.), *Proc. of ICCM'07*. Taylor and Francis, pp. 187-192.

[8] Cavazza, M., Charles, F., Mead, S. (2003). Interacting with virtual characters in interactive storytelling. In: Alonso, E., Kudenko, D., and Kazakov, D. (eds.), *Adaptive Agents and Multi-Agent Systems*. Springer LNAI, volume 2636, pp. 318–325.

[9] Damasio, A.R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Putnam.

[8] The assumptions that emotion elicitation can be described as a sudden increase or decrease on ERL and that emotion regulation is a smooth, long-term process which can be described by differential equations are quite common in the literature (e.g., [14]).

[10] Ekman, P., Friesen, W.V., and Ellsworth, P. (1972). *Emotion in the human face: guidelines for research and integration of Findings*. New York: Pergamon.

[11] Evans, J.S.B.T. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, vol. 7, pp. 454-459.

[12] Frijda, N.H., Manstead, A.S.R., and Bem, S. (eds.) (2000). *Emotions and Beliefs: How Feelings Influence Thoughts*. Cambridge Univ. Press, NY.

[13] Gärdenfors, P. (2001). Slicing the Theory of Mind. In: *Danish yearbook for philosophy*, 36, pp. 7-34. Museum Tusculanum Press.

[14] Gross, J.J. (1998). The Emerging Field of Emotion Regulation: An Integrative Review. *Review of General Psychology 2*, no. 3, pp. 271-299.

[15] Hoorn, J. F., Konijn, E. A., and Van der Veer, G. C. (2003). Virtual Reality: Do not augment realism, augment relevance. *Upgrade - Human-Computer Interaction: Overcoming Barriers*, vol. 4, issue 1, pp. 18-26.

[16] Jiang, H., Vidal, J.M., and Huhns, M.N. (2007). EBDI: An Architecture for Emotional Agents. In: *Proc. of the 6th Int. Joint Conf. on Autonomous Agents and Multi-Agent Systems, AAMAS'07*. ACM Press, pp. 38-40.

[17] Marsella, S., and Gratch, J. (2003). Modeling coping behavior in virtual humans: Don't worry, be happy. In: *Proceedings of Second International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS'03*. ACM Press, 2003, pp. 313-320.

[18] Marsella, S.C., Pynadath, D.V., and Read, S.J. (2004). PsychSim: Agent-based modeling of social interaction and influence. In: Lovett, M. et al. (eds.), *Proc. of ICCM'04*, pp. 243-248 Pittsburg, Pensylvania, USA.

[19] Pelachaud, C., Martin, J.C., Andre, E., Chollet, G., Karpouzis, K., and Pele, D. (eds.) (2007). *Intelligent Virtual Agents. Proc. of the 7th Int. Conf. on Intelligent Virtual Agents, IVA'07*. Springer LNAI, vol. 4722.

[20] Pereira, D., Oliveira, E., and Moreira, N. (2006). Modelling Emotional BDI Agents. In: *ECAI Workshop on Formal Approaches to Multi-Agent Systems (FAMAS 2006)*, Riva del Garda, Italy.

[21] Rao, A.S. and Georgeff, M.P. (1991). Modelling Rational Agents within a BDI-architecture. In: Allen, J., et al.. (eds.), *2nd Int. Conf. on Principles of KR. and Reasoning, KR'91*. Morgan Kaufmann, pp. 473-484.

[22] Silverman, B.G. (2001). More Realistic Human Behavior Models for Agents in Virtual Worlds: Emotion, Stress, and Value Ontologies. T.R., Univ. of Penn./ACASA.

[23] http://www.few.vu.nl/~tbosse/virtualagents/