

Incorporating Emotion Regulation into Virtual Stories

Tibor Bosse, Matthijs Pontier, Ghazanfar F. Siddiqui, and Jan Treur

Vrije Universiteit Amsterdam, Department of Artificial Intelligence,
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
{tbosse, mpontier, ghazanfa, treur}@few.vu.nl
<http://www.few.vu.nl/~{tbosse, mpontier, ghazanfa, treur}>

Abstract. This paper presents an approach to incorporate emotion regulation as addressed within psychology literature into virtual characters. To this end, first Gross' informal theory of emotion regulation has been formalised using a dynamical system style modelling approach. Next, a virtual environment has been created, involving a number of virtual agents, which have been equipped with the formalised model for emotion regulation. This environment has been used to successfully generate a number of emergent virtual stories, in which characters regulate their emotions by applying regulation strategies such as situation selection and attentional deployment. The behaviours shown in the stories were found consistent with descriptions of human regulation processes.

1 Introduction

In recent years, there has been an increasing interest in the area of *virtual storytelling*, addressing the development of computer systems that generate fictive stories in which the characters show realistic behaviour. In order to develop virtual stories, a large variety of approaches have been proposed, e.g., [4], [5], [13]. A trend that can be observed in many of these approaches is the movement from stories with a fixed, pre-scripted storyline towards *emergent narrative*, i.e., stories in which only a number of characters and their personalities are fixed, rather than the precise script of the story [1]. In the latter type of storytelling, ideally, all the designer (or writer) has to do is to determine which (types of) characters will occur in the play (although usually it is still needed to roughly prescribe the course of events). Hence, advantages of emergent narrative are the reduced amount of work that has to be spent by the writer, and the non-deterministic and unpredictable behaviour of the story.

In parallel with the shift from fixed storylines to emergent narrative, there has been a development in the nature of the involved characters as well. Recently, the characters (or agents) that are present in virtual stories are transforming more and more from shallow avatars to complex personalities with human-like properties such as emotions and theories of mind, e.g., [15]. To accomplish this, researchers have started to incorporate cognitive models within virtual characters, e.g., [10], [12]. Despite these first promising attempts, there is still a wide area to explore when it comes to enhancing virtual agents with cognitive capabilities.

In line with the development described above, this paper explores the possibilities to equip the characters involved in virtual stories with the capability of *emotion regulation*.

Informally, emotion regulation can be described as the process humans undertake to increase, maintain or decrease their emotional response, see e.g., [7], [8], [11], [14]. The idea is that, by offering virtual agents the capacity to actively regulate their emotions, they will be able to select those kinds of behaviours that they feel most comfortable with. As a result, such agents will 1) behave more realistically and 2) have more freedom in the choice of their actions, which enhances the emergent narrative effect. This approach is similar to the approach taken in [9], which aims at incorporating coping behaviour into virtual humans.

In order to build emotion regulation into virtual stories, in this paper the informal model by Gross [7] as found in psychology literature was taken as a basis. This model describes a number of strategies humans use to adapt their emotion response levels, varying from situation selection to cognitive change and response modulation. Next, this model has been formalised using a dynamical system style modelling approach (see also [3] for some initial steps). In addition, a virtual environment has been created, incorporating a number of virtual agents, and these agents have been equipped with the formalised model for emotion regulation. To test the behaviour of the model, a series of simulation experiments has been performed using the LEADSTO simulation language [2]. The model has been connected to the Vizard Virtual Reality Toolkit [16], to visualise the resulting stories in a graphical environment.

2 Emotion Regulation in the Virtual Agent Context

Gross [8] describes a process model of emotion regulation using the following definition: ‘Emotion regulation includes all of the conscious and nonconscious strategies we use to increase, maintain, or decrease one or more components of an emotional response’. In his model, Gross distinguishes four different types of emotion regulation strategies, which can be applied at different points in the process of emotion generation. First of all, when applying *situation selection*, a person chooses to be in a situation that matches the emotional response level the person wants to have for a certain emotion. For example, you can stay home instead of going to a party, because you are in conflict with someone who is going to that party. Second, when applying *situation modification*, a person modifies an existing situation so as to obtain a different level of emotion. For instance, when watching an irritating television program, you zap to another channel. Third, *attentional deployment* refers to shifting your attention to a certain aspect. For example, you close your eyes when watching an exciting penalty shoot-out. Finally, *cognitive change* refers to selecting a cognitive meaning to an event. For example, when a person loses a tennis match and blames the weather circumstances, instead of his own capacities.

To incorporate these strategies into virtual characters, a modelling approach was used that is based on the LEADSTO simulation environment [2] and the Vizard Virtual Reality Toolkit [16]. Due to space limitations, the technical details of LEADSTO and Vizard are not shown here. However, they can be found in Appendix A in [17]. Below, in Section 2.1, at a language-independent level a global overview is given of the model, of which an initial version can be found in [3]. Next, in Section 2.2, for each of the regulation strategies it is shown how it is used in the virtual agents playing as characters in virtual stories. The complete formal specification of the model (in LEADSTO notation) is shown in Appendix B in [17].

2.1 Global Overview

In order to incorporate emotion regulation strategies into virtual agents, a virtual environment is created that is populated by a number of agents. Each agent is equipped with a mechanism to regulate its emotions, which is based on the model as described informally by Gross [7]. To create a formal model, for any given type of emotion a number of variables have been introduced. For convenience, the model concentrates on one specific type of emotion. In principle, this can (at least) be any emotion that is considered to be a basic human emotion, e.g., sadness, happiness, or anger [6]. In order to describe the regulation of such an emotion, the model takes into account the four strategies discussed by Gross are used (i.e., *situation selection*, *situation modification*, *attentional deployment*, and *cognitive change*). Based on the four strategies mentioned, in the formalisation four corresponding *elements* (denoted by k) are introduced, for the objects that are affected by the particular strategies: *situation*, *sub-situation*, *aspect*, and *meaning*.

The model assumes that each agent aims at an optimal level of emotion. The regulation process in the virtual agents starts by comparing the actual *emotion response level* ERL to the emotion response level ERL_norm aimed at. The difference between the two is the basis for adjustment of the choices made for each of the elements k ; based on these adjusted choices, each element k will provide an adjusted *emotional value* EV_k .

To obtain a quantitative model, the emotion response level and the emotional values for the different elements for a given type of emotion are represented by real numbers in the interval $[0, 2]$ (where 0 is the lowest possible ERL (e.g., extreme sadness), and 2 the highest (e.g., extreme happiness)). In the model, the level of emotion to aim at (the ERL norm), is also expressed in a real number in the domain $[0, 2]$. Based on these concepts, the ERL is recalculated each step by the following difference equation formula:

$$\text{new_ERL} = (1-\beta) * \sum_k (w_k * EV_k) + \beta * \text{ERL}$$

In this formula, new_ERL is the new emotion response level, and ERL is the old emotion response level. The persistency factor β is the proportion of the old emotion response level that is taken into account to determine the new emotion response level. Initial tests have indicated that values for β around 0.9 deliver realistic results. The new contribution to the emotion response level is calculated by the weighted sum of the emotional values: $\sum_k w_k * EV_k$. By normalisation, the sum of all the weights w_k is taken to be 1. The following section describes how the different strategies influence the values of EV_k .

2.2 Emotion Regulation Strategies

Situation selection: which agent to meet. Every step, each agent chooses to be alone, or to contact another agent, by comparing the EVs it attaches to being alone and to being with other agents. The agent will always choose the option with the EV that is closest to its optimal ERL. When two agents contact each other, they decide to meet. When the agents are meeting, their EV for situation is set to the EV they attach to the other agent. When an agent chooses to be alone, its EV for situation is set to its EV for being alone.

Situation modification: what to talk about. When two agents are in a meeting, they will talk about a certain conversation subject. To decide which of the agents will start talking, each agent has a personal *dominance factor*. The agent with the highest dominance factor will choose the first conversation subject. Each step after this, the agent who has not chosen the current conversation subject will choose the next conversation subject.

When an agent gets to choose which conversation subject to talk about, it will compare the EVs it attaches to each conversation subject, and select the one that is closest to its optimal ERL. The EV for subsituation is set to the EV the agents attach to the conversation subject they are currently talking about. When an agent is not in a meeting, its EV for subsituation will be set to the neutral value of 1, since the agent is not in a subsituation. When an agent A talks to another agent B about a certain conversation subject CS, this will affect the way agent B thinks about agent A. Agent B's EV for agent A will change using the following formula:

$$\text{new_EV}_{\text{agent_A}} = \beta_{\text{friendship}} * \text{EV}_{\text{agent_A}} + (1 - \beta_{\text{friendship}}) * \text{EV}_{\text{CS}}$$

In this formula, $\text{new_EV}_{\text{agent_A}}$ is the new EV agent B will attach to agent A and $\text{EV}_{\text{agent_A}}$ is the old EV agent B attached to agent A. The persistency factor $\beta_{\text{friendship}}$ is the proportion of the old EV that is taken into account to determine the new EV. Here, values for $\beta_{\text{friendship}}$ bigger than 0.9 (where $\beta_{\text{friendship}}$ will get bigger when an agent knows another agent for a longer time) deliver realistic results. The new contribution to the ERL is determined by EV_{CS} : the EV agent B attaches to the conversation subject agent A is talking about. So how much an agent likes another agent, depends on how much an agent likes the conversation subjects another agent talks about.

The extent to which an agent likes to talk about a certain conversation subject can be changed by external events. For example, an agent will start to like a sports team more when this team wins a match. To accomplish this, the following formulas are used:

$$\text{new_EV}_{\text{CSn}} = \text{EV}_{\text{CSn}} + \Delta \text{EV}_{\text{CSn}}$$

$$\text{When a positive event occurs: } \Delta \text{EV}_{\text{CSn}} = \eta * \text{EV}_{\text{CSn}} * (d_{\text{max}} - \text{EV}_{\text{CSn}})$$

$$\text{When a negative event occurs: } \Delta \text{EV}_{\text{CSn}} = -\eta * \text{EV}_{\text{CSn}} * (d_{\text{max}} - \text{EV}_{\text{CSn}})$$

In these formulas, $\text{new_EV}_{\text{CSn}}$ is the new EV the agent attaches to the conversation subject, and EV_{CSn} is the old EV the agent attached to the conversation subject. Here η is a variable that determines the speed of adjusting EVs to conversation subjects. A lower η will result in slower adjustment. Here, an η of 0.02 delivers realistic results. The part $\text{EV}_{\text{CSn}} * (d_{\text{max}} - \text{EV}_{\text{CSn}})$ prevents EV_{CSn} from under- or overadjustment.

Attentional deployment: on which aspect to focus. When an agent is in a conversation, it can choose to pay attention to, or to distract its attention from the conversation. Every step, the agent chooses the option with the EV closest to its optimal ERL. The EVs the agent attaches to paying attention or distracting its attention, depend on the conversation subject the agent is currently talking about, according to the following formulas:

$$\text{new_EV}_{\text{pay_attention}} = \beta_{\text{asp}} * \text{EV}_{\text{pay_attention}} + (1 - \beta_{\text{asp}}) * \text{EV}_{\text{CS}}$$

$$\text{new_EV}_{\text{distract}} = \beta_{\text{asp}} * \text{EV}_{\text{distract}} + (1 - \beta_{\text{asp}}) * (-\text{EV}_{\text{CS}} + d_{\text{max}})$$

In these formulas, $\text{new_EV}_{\text{pay_attention}}$ and $\text{new_EV}_{\text{distract}}$ are the new EVs for pay_attention and distract, and $\text{EV}_{\text{pay_attention}}$ and $\text{EV}_{\text{distract}}$ are the old EVs for pay_attention and distract. The persistency factor β_{asp} is the proportion of the old EV that is taken into account to determine the new EV. The new contribution to the EV for pay_attention is determined by EV_{CS} , the EV the agent attaches to the conversation subject it is talking about. The new contribution to the EV for distract is calculated by $(-\text{EV}_{\text{CS}} + d_{\text{max}})$. This will reach a high value when the agent attaches a low EV to the conversation subject, and a low value when the agent attaches a high value to the conversation subject. So when the agent likes the conversation subject, it will be more likely to pay attention to the conversation. The agent chooses to distract from, or pay attention to the conversation, by comparing the two EVs for paying attention and distracting, and picking the option with the EV closest to its optimal ERL.

Cognitive change: which meaning to attach. Every step, agents can choose to apply self-talk. An agent can use self-talk to relativise its current state of mind, or on the other hand, to attach more meaning to its current state. Every step, an agent chooses to relativise, attach a stronger meaning, or to apply no self-talk, by picking the option with the EV closest to the optimal ERL of the agent. The EV for not applying self-talk always has the neutral value of 1. The EVs for relativising and attaching more meaning depend on the ERL of the agent, and are updated every step according to the following formula's:

$$\begin{aligned} \text{new_EV}_{\text{relativise}} &= d_{\text{max}} - \text{ERL} \\ \text{new_EV}_{\text{attach_more_meaning}} &= \text{ERL} + (\text{ERL} - 1) * (1 - \text{abs}(1 - \text{ERL})) \end{aligned}$$

When an agent has a high ERL, the EV for relativising will be low, and when an agent has a high ERL, the EV for attaching more meaning will be high. So relativising always influences the ERL of the agent to reach a more neutral value.

When the ERL of the agent has the neutral value of 1, $(\text{ERL} - 1)$ will be 0, and the EV for attaching more meaning will be 1. When the ERL of the agent is smaller than 1, then $\text{ERL} - 1$ will have a negative value, and the EV for attaching more meaning will have a value that is smaller than the current ERL. When the ERL of the agent is bigger than 1, then $\text{ERL} - 1$ will be bigger than 1, and the EV for attaching more meaning will have a value that is bigger than the current ERL. So attaching more meaning always influences the ERL of the agent to a more extreme value than the current one. Multiplying by $(1 - \text{abs}(1 - \text{ERL}))$ prevents the EV from reaching values that are out of the domain.

3 Simulation Experiments

Several experiments have been done to test the simulation model's ability to generate interesting scenarios. To obtain movies in Vizard, events in the LEADSTO simulations were translated to visualisations in Vizard. The exact mapping that was used for this translation is shown in Appendix C in [17]. For example, the fact that an agent is happy is visualised by a certain type of smile, and the fact that an agent distracts from a conversation is visualised by this agent moving its head away from its conversation partner.

In all of the simulations, three agents are involved, which will be called Barry, Gary, and Harry. The particular emotion these agents will try to regulate during the scenario's is their amount of happiness. To enable this, the particular topics they are allowed to talk about are football (in particular, the Dutch football teams Ajax and Feyenoord) and hockey. The parameter settings of all agents used in three specific experiments are shown in Appendix D in [17].

Due to space limitations, only one of the simulation experiments is discussed in this paper. The results of the LEADSTO simulation of this experiment can be seen in Figure 1. Here, time is on the horizontal axis, whereas different events are displayed on the vertical axis. A dark box on top of a line indicates that an event is true at that time point; a light box below a line indicates that an event is false. A detailed description of what happens in this scenario is provided in Appendix E in [17].

As mentioned earlier, using a specific conversion program that has been implemented, LEADSTO simulations were translated into movies in Vizard. A screenshot of an example Vizard movie (which corresponds to the scenario shown in Figure 1) is shown in Figure 2. This figure shows a situation in which (on the foreground) two agents are having a conversation. The left agent is talking about hockey, but the right agent tries to

distract from the conversation by moving its head away from the conversation. The cognitive meaning that each agent attaches to its current thoughts is displayed (in red) above the heads of the agents. Meanwhile, in the background a third agent is standing alone. The full Vizard movie of this scenario (as well as the movies that correspond to the two other experiments described in Appendix D) can be found on [17].

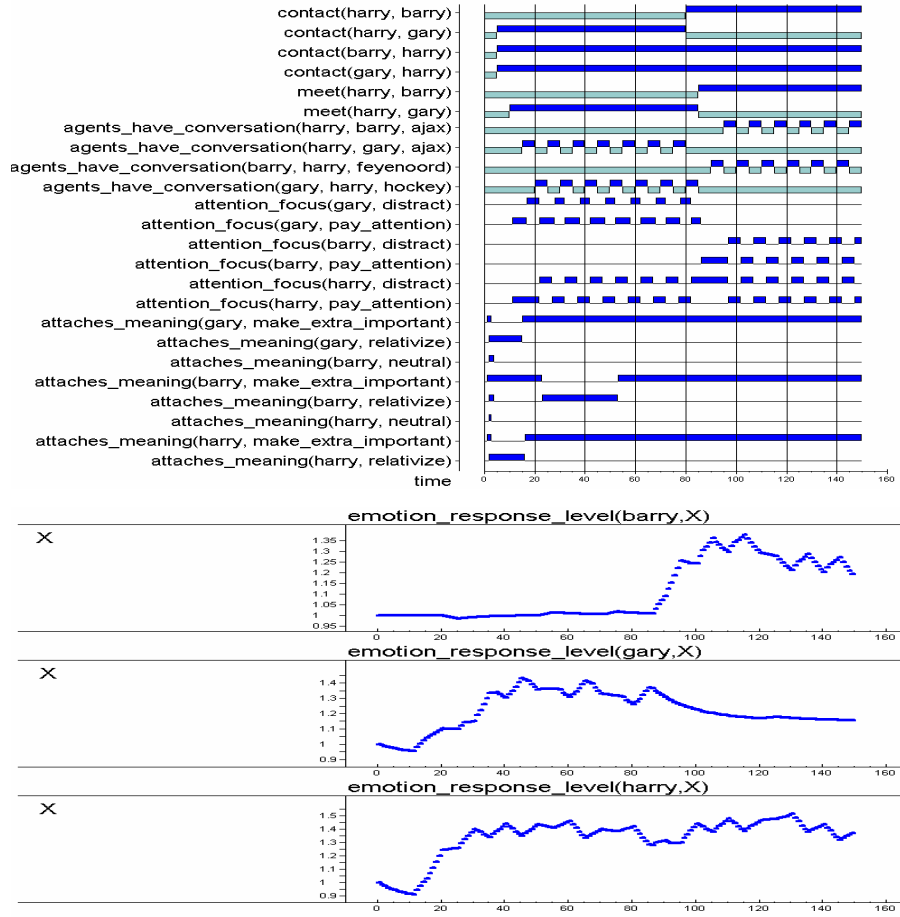


Fig. 1. Example Simulation Trace

The resulting movies provide a first indication that the emotion regulation strategies as described by [7] have been implemented successfully within the virtual agents used as characters. To be specific, the agents are able to perform *situation selection* by selecting different conversation partners, and withdrawing from conversations if desired. Moreover, they can perform *situation modification* by changing conversation topics, they can perform *attentional deployment* by changing the amount of attention they pay to a conversation, and they can perform *cognitive change* by changing the cognitive meaning they assign to their thoughts (e.g., by stating to themselves that something is not very important). These behaviours were found consistent with predicted behaviours for situations as described by Gross [7], [8] (which are based on empirical evidence).



Fig. 2. Screenshot of an example scenario in Vizard

4 Discussion

Within the domain of virtual storytelling, the idea of *emergent narrative* has become more and more popular [1]. Moreover, there is a growing trend to incorporate cognitive models within the characters involved in virtual stories (e.g., [10], [12]). As a next step in that direction, the current paper aims at building emotion regulation as known from psychology literature into virtual characters. To this end, the informal model by Gross [7] was taken as a basis, and has been formalised using a dynamical system style modelling approach (see also [3] for some initial steps). A virtual environment has been created, which includes a number of virtual agents that have been equipped with the formalised model for emotion regulation. To test the behaviour of the model in a prototyping phase, a series of simulation experiments has been performed using the LEADSTO simulation language [2]; in the Vizard Virtual Reality Toolkit [16], such simulations have been visualised in a graphical environment. The resulting movies provide a first indication that the emotion regulation strategies as described by [7] have been implemented successfully within the virtual characters. The simulation results have been compared with the behaviours for different situations as described by Gross [7], [8], and found consistent. Validation involving comparison with detailed empirical data is left for future work.

Concerning related work, an approach in the literature that has similarities to the current approach is [9]. In that paper, a computational model is introduced that can simulate several strategies about how humans cope with emotions, such as ‘positive reinterpretation’ and ‘denial’. Their approach makes use of plan-based causal representations, augmented with decision-theoretic planning techniques, whereas our approach uses dynamical systems representations. Other differences are that they propose a “content model”, in which appraisal and regulation operate on rich representations of the emotion-evoking situation, and that their work has been evaluated against clinical data.

Virtual stories involving characters with elaborated cognitive or psychological capabilities can be used for a number of purposes. On the one hand, they may be used for entertainment (e.g., for creating computer games with more complex, unpredictable and

more human-like characters). On the other hand, they may be used for educational purposes (e.g., to create a virtual training environment for psychotherapists, which enables them to practice anger management sessions with virtual clients). Further research will investigate whether the model is suitable for such purposes. As soon as these types of challenges will be tackled, also a more precise evaluation will be performed of how humans perceive the current characters (e.g. in terms of believability).

References

1. Aylett, R. (1999). Narrative in Virtual Environments: Towards Emergent Narrative. In: *Working notes of the Narrative Intelligence Symposium*, AAAI Spring Symposium Series. Menlo Park, California, AAAI Press.
2. 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*. To appear, 2007.
3. Bosse, T., Pontier, M., and Treur, J. (2007). A Dynamical System Modelling Approach to Gross' Model of Emotion Regulation. In: *Proceedings of the 8th International Conference on Cognitive Modeling, ICCM'07*. Taylor and Francis, to appear.
4. Cavazza, M., Charles, F., and Mead, S. (2002). Interacting with virtual characters in interactive storytelling. In: *Proceedings of the First International Conference on Autonomous Agents and Multi-agent Systems, AAMAS'02*. ACM Press, pp. 318-325.
5. Dautenhahn, K. (1998). Story-Telling in Virtual Environments. *Proceedings of the ECAI'98 Workshop on Intelligent Virtual Environments*, Brighton, UK.
6. Ekman, P., Friesen, W.V., and Ellsworth, P. (1972). *Emotion in the human face: guidelines for research and integration of Findings*. New York: Pergamon.
7. Gross, J.J. (1998). The Emerging Field of Emotion Regulation: An Integrative Review. *Review of General Psychology*, vol. 2, no. 3, pp. 271-299.
8. Gross, J.J. (2001). Emotion Regulation in Adulthood: Timing is Everything. *Current directions in psychological science*, vol. 10, no. 6, pp. 214-219.
9. 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, pp. 313-320.
10. Marsella, S. C., Johnson, W. L., and LaBore, C. (2000). Interactive Pedagogical Drama. In: *Proceedings of the 4th International Conf. on Autonomous Agents*, ACM Press, pp. 301-308.
11. Ochsner, K.N., and Gross, J.J. (2005). The cognitive control of emotion. *Trends in Cognitive Sciences*, vol. 9, pp. 242-249.
12. Paiva, A., Machado, I., and Prada, R. (2001). Heroes, villains, magicians, ...: Dramatis personae in a virtual story creation environment. In: *Proceedings of the Conference on Intelligent User Interfaces, IUI'01*, pp. 129-136.
13. Theune, M., Faas, S., Heylen D., and Nijholt, A. (2003). The Virtual Storyteller: Story Creation by Intelligent Agents. *Proceedings of Technologies for Interactive Digital Storytelling and Entertainment*, pp. 204-215.
14. Thompson, R.A. (1994). Emotion regulation: A theme in search of definition. In N.A. Fox (ed.), *The development of emotion regulation: Biological and behavioral aspects. Monographs of the Society for Research in Child Development*, vol. 59, pp. 25-52.
15. Van Vugt, H.C., Hoorn, J.F., Konijn, E.A., and De Bie Dimitriadou, A. (2006). Affective affordances: Improving interface character engagement through interaction. *International Journal of Human-Computer Studies*, vol. 64, no. 9, pp. 874-888.
16. Worldviz. *Vizard Virtual Reality Toolkit*. University of California. URL: <http://www.worldviz.com/vizard.htm>.
17. <http://www.cs.vu.nl/~tbosse/virtualstories>.