A Dynamical System Modelling Approach to Gross' Model of Emotion Regulation

Tibor Bosse (tbosse@few.vu.nl)   **Matthijs Pontier (mpontier@few.vu.nl)**   **Jan Treur (treur@few.vu.nl)**

Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081, NL-1081 HV, Amsterdam, The Netherlands

Abstract

This paper introduces a computational model for emotion regulation formalising the model informally described by Gross (1998). The model has been constructed using a high-level modelling language, and integrates both quantitative aspects (such as levels of emotional response) and qualitative aspects (such as decisions to regulate one's emotion). A number of simulation experiments have been performed, demonstrating that the computational model successfully reflects the model as described by Gross.

Introduction

Emotions were historically seen as neural activation states without a function. However, recent research provides evidence that emotions are functional (e.g., Damasio, 2000). Emotions have a facilitating function in decision making, prepare a person for rapid motor responses, and provide information regarding the ongoing match between organism and environment. Emotions also have a social function. They provide us information about others' behavioural intentions, and script our social behaviour (Gross, 1998). In the past two decades, psychological research has started to focus more on *emotion regulation* (e.g., Gross, 1998, 2001; Ochsner and Gross, 2005; Thompson, 1994). In brief, emotion regulation is the process humans undertake in order to affect their emotional response. Recent neurological findings (such as bidirectional links between limbic centers, which generate emotion, and cortical centers, which regulate emotion) have changed the consensus that emotion regulation is a simple, top-down controlled process (Gross, 1998).

This article introduces a computational model to simulate emotion regulation, based on the process model described informally by Gross (1998, 2001). Such a model can be used for different purposes. In the first place, from a Cognitive Science perspective, it can provide insight in the process of emotion regulation. This may be useful for the purpose of developing therapies for persons that have difficulties in regulating their emotions (Burns et al., 2003; Towl and Crighton, 1996), for example, in work with forensic inpatients. In addition, a model for emotion regulation can be used in the field of Artificial Intelligence, see e.g. (Bates, 1994). For example, in the domain of virtual reality it can be used to let virtual agents show human-like behaviour regarding emotion regulation. Finally, computational models for emotion regulation may play a role within the field of Ambient Intelligence (Aarts, Harwig, and Schuurmans, 2001). For instance, when humans have to interact intensively with automated systems, it is useful if the system maintains a model of the emotional state (and the emotion regulation process) of the user. This enables the system to adapt the interaction to the user's needs.

Below, first Gross's model of emotion regulation is briefly discussed. The 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, the dynamical system style modelling approach used is briefly introduced. After that, the simulation model formalising the model of Gross is described, and some simulation results are shown, both for ideal cases and for cases of over- and under-regulation. The paper concludes with a discussion.

Gross' Model for Emotion Regulation

Gross (2001) 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'

The components he considers are (1) the *experiential* component, (the subjective feeling of the emotion), (2) the *behavioural* component (behavioural responses), and (3) the *physiological* component (responses such as heart rate and respiration). Humans use strategies to affect their level of emotional response for a given type of emotion, for example, to prevent a person from having a too high emotional or too low emotional response level. He differentiates between antecedent-focused strategies and response-focused strategies. *Antecedent-focused strategies* are applied to the process preparing for response tendencies before they are fully activated. *Response-focused strategies* are applied to the activation of the actual emotional response, when an emotion is already underway.

In his model, Gross distinguishes four different types of antecedent-focused emotion regulation strategies, which can be applied at different points in the process of emotion generation: *situation selection, situation modification, attentional deployment* and *cognitive change*. A fifth strategy, *response modulation*, is a response-focused strategy. Figure 1 shows an overview of these strategies.

The first antecedent-focused emotion regulation strategy in the model is *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, a person can stay home instead of going to a party, because he is in conflict with someone who is going to that party. This is an example of down-regulating one's emotion.

The second antecedent-focused emotion regulation strategy in the model is *situation modification*. When this

strategy is applied, a person modifies an existing situation so as to obtain a different level of emotion. For instance, when watching an irritating television program, one may zap to another channel.
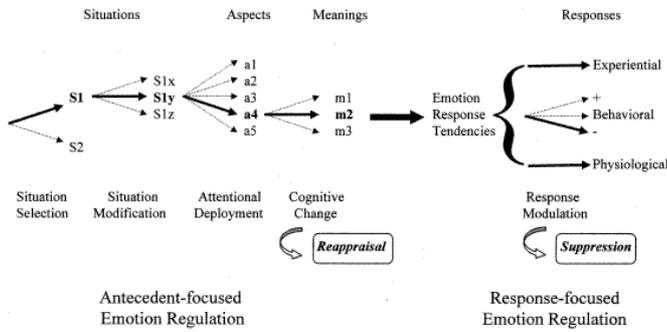

Figure 1: Emotion Regulation Model by Gross (1998).

The third antecedent-focused emotion regulation strategy is *attentional deployment*. This strategy refers to shifting your attention to a certain aspect. For example, one may close his eyes when watching an exciting penalty shoot-out. The fourth antecedent-focused emotion regulation strategy is *cognitive change*: selecting a cognitive meaning to an event. A specific type of cognitive change, which is aimed at down-regulating emotion, is *reappraisal*:

'Reappraisal means that the individual reappraises or cognitively re-evaluates a potentially emotion-eliciting situation in terms that decrease its emotional impact' (Gross, 2001).

An example of reappraisal is a case when a person loses a tennis match and blames the weather circumstances, instead of his own capacities. However, note that cognitive change could also be aimed at up-regulating emotion.

The fifth emotion regulation strategy, *response modulation*, a response-focused strategy, is applied after the emotion response tendencies have been generated: a person tries to affect the process of response tendencies becoming a behavioural response. A specific type of response modulation, again aimed at down-regulating, is *suppression*:

'Suppression means that an individual inhibits ongoing expressive behaviour.' (Gross, 2001).

An example of suppression is a person that hides being nervous when giving a presentation. As Gross considers response modulation to be not very effective, this strategy is not considered in the paper, although it would not be difficult to incorporate it in the computational model.

Modelling Approach

Modelling the various aspects involved in Gross' model in an integrated manner poses some challenges. On the one hand, qualitative aspects have to be addressed, such as decisions to regulate one's emotion (e.g., by selecting a different situation). On the other hand, quantitative aspects have to be addressed, such as levels of emotional response.

The modelling approach based on the modelling language LEADSTO (Bosse, Jonker, Meij, and Treur, 2007) fulfils these desiderata. It integrates qualitative, logical aspects such as used in approaches based on temporal logic (e.g., Barringer et al., 1996) with quantitative, numerical aspects such as used in Dynamical Systems Theory (e.g., Ashby, 1960; Port and van Gelder, 1995). Direct temporal dependencies between two state properties in successive states are modelled by *executable dynamic properties* defined as follows. Let a and b be state properties of the form 'conjunction of ground atoms or negations of ground atoms', then the notation $a \rightarrow\!\!\!\rightarrow_{e, f, g, h} b$ means:

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

Atomic state properties can have a qualitative, logical format (e.g., desire(d), expressing that desire d occurs), or a quantitative, numerical format (e.g., has_value(x, v) expressing that variable x has value v).

Global Overview of the Model

Gross has described his process model for emotion regulation informally. In order to be able to formalise his 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 be any emotion that is considered to be a basic human emotion, e.g., sadness, happiness, or anger (Ekman, Friesen, and Ellsworth, 1972).

In order to describe the regulation of such an emotion, the model takes into account a number of emotion regulation *strategies* that can be chosen. In the variant of the model as described in this paper, the four antecedent-focused emotion regulation strategies discussed by Gross are used (i.e., situation selection, situation modification, attentional deployment, and cognitive change). For the moment, response modulation is not considered. However, the model is generic in the sense that this set of strategies considered can easily be adapted. Based on the four strategies mentioned, in the formalisation four corresponding *elements* k are introduced, denoting the objects that are influenced by the particular strategies (see Table 1).

Table 1: Strategies and elements addressed in the model.

Strategy	Corresponding Element
situation selection	situation
situation modification	sub_situation
attentional deployment	aspect
cognitive change	meaning

In the model it is assumed that at each point in time, for each element k a certain choice is in effect, and this choice has a certain *emotional value* v_k attached. This emotional value contributes to the *emotion response level* ERL via an element-specific weight factor w_k, thereby taking into account a persistency factor β indicating the degree of persistence or slowness of adjusting of the emotion response

level when new emotional values are obtained. Someone whose emotions can change rapidly (e.g., who stops being angry in a few minutes after a fight) will have a low β.

Humans are always searching for a certain level of emotion depending on the person[1]. For instance, some enjoy extreme sports, while others prefer a more quiet kind of recreation. The level of emotion aimed at depends also on the type of emotion. Most humans aim at a relatively high level of emotion for happiness, while they aim at a lower level of emotion for fear. The regulation process starts by comparing the actual emotion response level ERL to the emotion response level ERL_norm aimed at. The difference d 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 have an adjusted emotional value v_k. The strength of such an adjustment is expressed by a modification factor α_k, which can be seen as a flexibility or willingness (conscious or unconscious) to change one's emotional value for a certain element. For instance, the α for the element 'situation selection' can be seen as the flexibility to change one's situation. An overview of the variables used in the model is given in Table 2.

Table 2: Variables addressed in the model.

Variable	Meaning
ERL	Emotion Response Level
ERL_norm	Emotion Response Level aimed at
D	Difference between ERL and ERL_norm
β	Persistency factor for ERL
K	Elements indicating strategies incorporated
w_k	Weight of element k in adjusting the ERL
v_k	Emotional value for element k
α_k	Modification factor that represents the flexibility to change the emotional value of element k

Some of these variables were chosen to be set at forehand and remain constant during the process (in particular ERL_norm, β, w_k, α_k). The other variables depend on each other and on the fixed variables, as shown in a qualitative manner in the graph depicted in Figure 2.
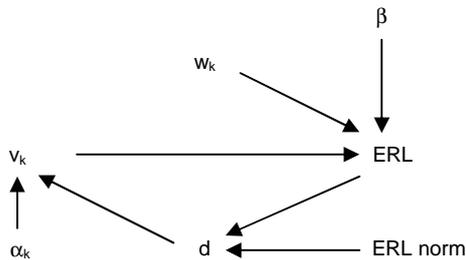

Figure 2: Dependencies between the variables.

This graph shows that the emotion response level ERL is affected by the emotional values v_k for the different elements, the weights w_k attached to these elements, and the persistency factor β that indicates in how far the previous response level affects the current one. The difference d between response level and norm obviously depends on both of these factors. Finally, the emotional values v_k for the different elements are affected by this difference d and the modification factor α_k.

The Quantitative Relations in the Model

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 emotion response level, and 2 the highest). In the model, a fixed level of emotion to aim at is assumed (the ERL norm), also expressed in a real number in the domain [0, 2]. As a simple illustration, suppose one wants to influence its state of anger by selecting an appropriate situation, and one deliberates whether to go to a party or not. This can be represented by introducing two different situations sit1 and sit2, for example with $v_{sit1}=1.5$ (since going to the party will increase the state of anger) and $v_{sit2}=0.5$ (staying home will decrease the state of anger). Moreover, the ERL norm can for instance be 0.7 (i.e., one aims at being a bit angry, but not too angry). In that case, if one's current ERL is already high, one will be likely to stay home (i.e., choose sit2), and vice versa.

The process of emotion regulation has a continuous nature. At any point in time, the characteristics of the current situation affect a person's emotional response level. Meanwhile, this emotional response level affects the person's choice for the emotional values v_k, which in turn influence the current situation (see also the cycle in Figure 2). An approach to model such a process is the Dynamical Systems Theory (DST) based on differential equations; e.g., (Port and van Gelder, 1995). To use differential equations for simulation, some form of *discretisation* is needed. Therefore, instead of differential equations, a set of difference equations is used, with a fixed step size s, that can be taken any size as desired.

Updating the Emotional Response Level

Based on the above ideas, the emotion response level is recalculated each step by the following difference equation formula:

$$\text{new\_ERL} = (1-\beta) * \Sigma_k (w_k * v_k) + \beta * \text{ERL}$$

In this formula[2], 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. The new contribution to the emotion response level is calculated by the weighted sum of the emotional values: $\Sigma_k w_k * v_k$. By normalisation, the sum of all the weights w_k is taken to be 1. According to the indication

---

[1] Although we use words like 'searching for' to describe this process, it is not claimed that this process is always a conscious, deliberate activity.

[2] Note that the formula can also be rewritten into the following difference equation format:

$\Delta\text{ERL} = (1-\beta) * (\Sigma_k (w_k * v_k) - \text{ERL}) \Delta t$ with $\Delta\text{ERL} = \text{new\_ERL} - \text{ERL}$

This format shows more explicitly how β determines the speed of adaptation of ERL to the new contribution $\Sigma_k w_k * v_k$; here Δt is taken 1.

of Gross (2001), elements that are affected at an earlier point in the emotion regulation process have higher weights. Within the simulation model, the update of the emotional response level is expressed by the following dynamic property in LEADSTO format (where s is the step size):

LP1 (Update Emotion Response Level)
emotion_response_level(erl)
and has_weight(situation, w1)
and has_weight(sub_situation, w2)
and has_weight(aspect, w3)
and has_weight(meaning, w4)
and has_emotional_value(situation, v1)
and has_emotional_value(sub_situation, v2)
and has_emotional_value(aspect, v3)
and has_emotional_value(meaning, v4)
$\rightarrow\!\!\rightarrow_{0, 0, s, s}$ emotion_response_level((1-beta) *
(w1*v1 + w2*v2 + w3*v3 + w4*v4) + beta * erl)

Updating the Emotional Values

The chosen emotional values v_k, which affect the emotion response level, are on their turn recalculated each step by the following set of difference equations:

$$d = ERL - ERL_{norm}$$
$$\Delta v_k = -\alpha_k * d / d_{max} \, \Delta t$$
$$new\_v_k = v_k + \Delta v_k$$

In these formulas, new_v_k is the new emotional value v_k, and old_v_k is the old emotional value v_k, while Δv_k is the change of the emotional value v_k (either positive or negative), and Δt the time step, which is taken 1 in this paper. The change in the emotional value v_k is calculated by the formula $-\alpha_k * d / d_{max}$. In this formula, α_k is the modification factor, and d is the difference between the actual emotion response level and the desired emotion response level (represented by ERL_norm). Here d_{max} is an estimation of the maximum difference that can be reached. So d / d_{max} is the proportion of the maximal reachable level of emotion above the level of emotion aimed at (or below this level, if d is negative).

When the actual emotion response level equals the desired emotion response level, then $d = 0$; this means that $\Delta v_k = 0$, so the emotion response level will not change. Moreover, a person will 'choose' an element with a more extreme emotional value v_k when (s)he is more flexible in this emotional value v_k (this is the case when α_k is high), or when (s)he experiences an emotion response level that is further away from the desired emotion response level (this is the case when d deviates more from 0). Within the simulation model, the update of emotional values is expressed as follows:

LP2 (Update Emotional Values)
emotion_response_level(erl) and erl_norm(erl_norm)
and has_emotional_value(element, v)
and has_modification_factor(element, a)
$\rightarrow\!\!\rightarrow_{0, 0, s, s}$
has_emotional_value(element, v – a * (erl - erl_norm) / dmax)

Simulation Results

A number of experiments have been performed to test what kind of behaviour can be simulated. Each subsection below addresses a specific type of scenario. Two types of cases are

addressed: those with an optimal form of regulation (compared to the emotion response level aimed at), and cases of over- and under-regulation. The different scenarios are established by taking different settings for the modification factors α_k. The values of the other variables are the same for all experiments described in this section, see Table 3.

Table 3: Values of variables used in the simulations.

Variable	Fixed Value		Variable	Initial Value
ERL_{norm}	0.5		ERL	1.85
β	0.7		v_1	1.90
w_1	0.35		v_2	1.85
w_2	0.30		v_3	1.80
w_3	0.20		v_4	1.75
w_4	0.15			
s	1			

As shown in the table, the person considered has an optimal level of emotion of 0.5 in the domain [0, 2]. The factor β is set to 0.7, which means that in each step, 70% of the old emotional response level persists, and the remaining 30% is determined by the new emotional values. The weight attached to situation selection is 0.35, which means that the selected situation determines 35% of the 30% of the new emotion response level that is determined by the emotional values. Similarly, the weights for situation modification, attentional deployment, and cognitive change are set to 0.30, 0.20, and 0.15, respectively. The results of the experiments are shown and explained below.

Optimal forms of emotion regulation

In the first experiment, all modification factors α_k were set to 0.15. The results are shown in Figure 3. In such figures, time is on the horizontal axis; the values of the different variables are shown on the vertical axis.
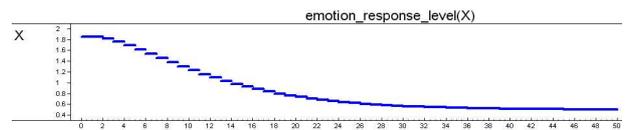

Figure 3: Results for an optimal case (equal α_k).

The emotional response level decreases monotonically without decreasing below the level aimed at. So, the subject gradually reaches his level of emotion aimed at. The emotional values show similar behaviour (due to space limitations not shown here).

In the second experiment, the subject has for each element k a different flexibility α_k in emotion regulation:
$$\alpha_1 = 0.20, \quad \alpha_2 = 0.15, \quad \alpha_3 = 0.10, \quad \alpha_4 = 0.05$$
The results of this experiment are shown in Figure 4. Here, the emotion response level reaches the emotion response level of 0.5 aimed for in a reasonable amount of time, just like in the optimal case. However, the way the emotional values change in order to achieve this differs from the first experiment. Here, it is important to note that

the scale on the vertical axis is not the same for the different graphs in Figure 4. The graphs show that the emotion response levels of the elements with a higher α descend much quicker and further than the elements with a lower α. For example, situation selection (α=0.20) has reached an emotional value of 0 at the end of the simulation, whereas cognitive change (α=0.01) changes only a little bit, and reaches an emotional value of about 1.3. This means that the subject finds a way to reach his/her level of emotion aimed for, and does this by changing his/her behaviour more for the elements for which (s)he has a higher flexibility.
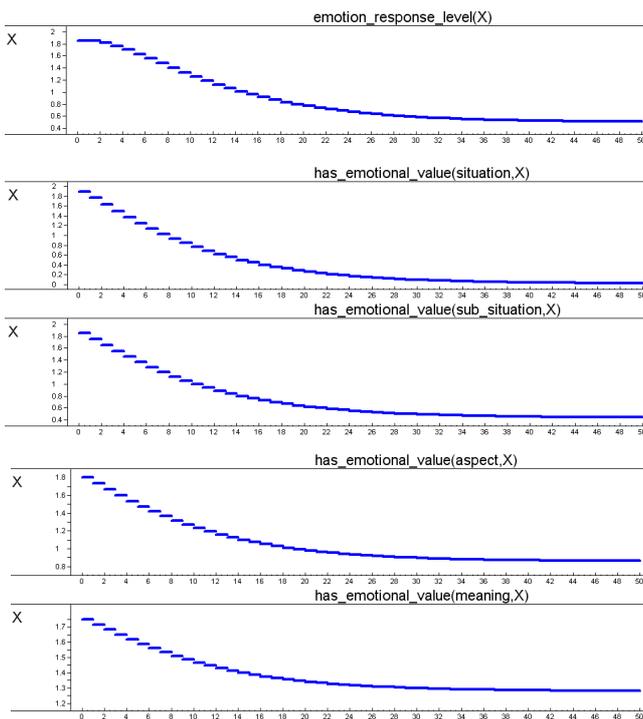

Figure 4: Results for an optimal case (different α_k).

Over- and under-regulation

In the third experiment, the modification factors α_k for all elements were set to 0.4. This means that the subject has a relatively high flexibility in emotion regulation, for all elements. The behaviour of the emotion response level in this experiment is shown in Figure 5.

In this case, the emotion response level starts to decrease rapidly, immediate after the experiment has started. However, it decreases below the level of 0.5 aimed at. It reaches its minimum after 15 steps in the simulation, at about 0.3: the subject over-regulates his/her emotion. After this, the emotion response level starts to raise until it is just above the optimal level of 0.5, and stays more or less at this value aimed at for the rest of this simulation.

The lower part of Figure 5 shows how the subject changed his/her emotional values in order to achieve this. These emotional values all show similar behaviour, since

the α_k's, which represent the flexibility and willingness to change behaviour, were set to the same value. Also, the graphs of the emotional values are comparable to the graph of the emotion response level. The emotional values make a somewhat steeper curve, especially at the start of the graph. This makes sense, because the emotion response level is only for 30% determined by the emotional values, and for 70% by its own old value.
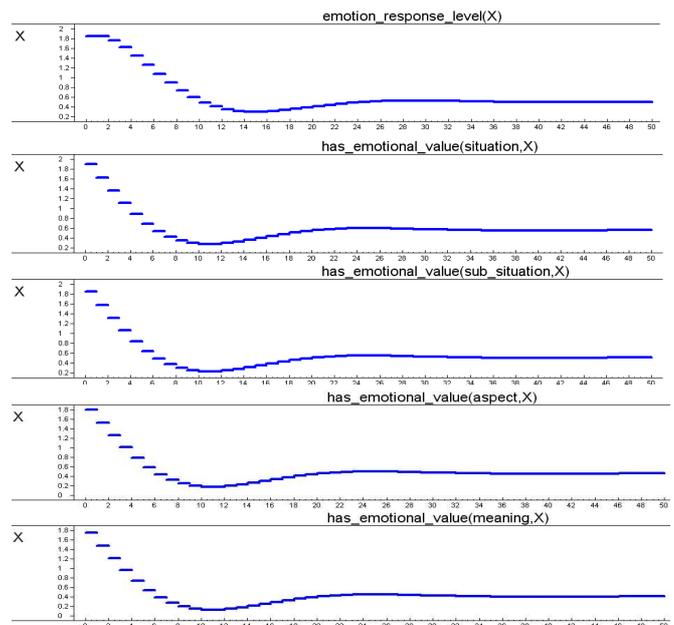

Figure 5: Results for the over-regulation case.

In the fourth experiment presented, the subject has a very low flexibility in emotion regulation, with an α_k value of 0.01 for all elements. The results of this experiment are shown in Figure 6. In this experiment, the emotion response level decreases extremely slowly: under-regulation. After 50 steps, it has only decreased by 0.3 until 1.55, as can be seen in the graph.
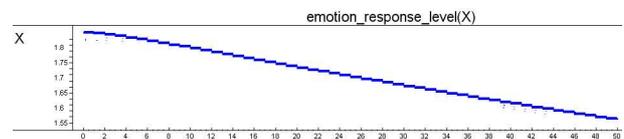

Figure 6: Results for the under-regulation case.

Discussion

In this paper, a formal model for Gross' (informally described) model of emotion regulation has been introduced. The emotion regulation model has been constructed using the high-level simulation language LEADSTO as a modelling vehicle, and integrates both quantitative, dynamical system aspects (such as levels of emotional response) and qualitative aspects (such as decisions to regulate one's emotion). Simulation

experiments have been performed for different situations, by using different settings for the modification factors α_k: for ideal cases (all α_k are medium, or the α_k have different values), for cases of over-regulation (all α_k are high), and for cases of under-regulation (all α_k are low). The experiments show that different values for the modification factors α_k indeed result in different patterns.

As a preliminary validation of the model, the simulation results have been compared with the predicted behaviours for different situations as described by Gross, which are (partly) based on empirical evidence (Gross, 1998, 2001). The patterns produced by the model were found consistent with Gross' descriptions of examples of human regulation processes. Validation involving extensive comparison with detailed empirical data is left for future work.

Although the process of emotion regulation is widely investigated in the literature (e.g., Gross, 1998, 2001; Ochsner and Gross, 2005; Thompson, 1994), not so many contributions address the possibility of developing a computational model of this process. The computational models that have been developed so far either address some very specific aspects of the process at a more detailed (neurological) level, see e.g. (Thayer and Lane, 2000), or they aim at incorporating emotions into software agents, in which case they focus more on emotion elicitation (appraisal) than on emotion regulation, see e.g. (Armony et al., 1997; Bates, 1994; Velasquez, 1997). The current paper can be seen as an attempt to build a bridge between both directions. It formalises an existing theory about emotion regulation using a high-level modelling language, but still in enough detail to be able to generate useful simulation traces. As such, it has similarities with the work by Marsella and Gratch (2003), who propose an approach to incorporate both appraisal and coping behaviour into virtual humans. Their approach makes use of plan-based causal representations, augmented with decision-theoretic planning techniques, whereas our approach uses dynamical systems represen-tations. 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.

The presented model is still in an early stage of development. For example, the modification factors α_k are currently fixed. In order to make the model adaptive, these factors can be made adjustable. A way to accomplish this is to adapt the values of the α_k to one's satisfaction about the past emotion regulation process. This way, the model could simulate cases in which humans learn to select the ideal situations, as in certain types of therapy. Another possible extension to the model would be to make the desired emotion response level ERL_norm dynamic, so that it can depend on specific circumstances. A final extension would be to represent the different elements k using more complex knowledge structures, and to enable the model to dynamically derive the different emotional values from these structures, as is done, for example, in (Marinier and Laid, 2004). Future work will explore such possibilities.

References

Aarts, E., Harwig, R., and Schuurmans, M. (2001). Ambient Intelligence. In *The Invisible Future: The Seamless Integration of Technology into Everyday Life*, McGraw-Hill, 2001.

Armony, J.L., Servan-Schreiber, D., Cohen, J.D., and Ledoux, J.E. 1997). Computational modeling of emotion: Explorations through the anatomy and physiology of fear conditioning. *Trends in Cognitive Sciences*, vol. 1, pp. 28-34.

Ashby, R. (1960). *Design for a Brain*. Second Edition. Chapman & Hall, London. First edition 1952.

Barringer, H., Fisher, M., Gabbay, D., Owens, R., and Reynolds, M. (1996). *The Imperative Future: Principles of Executable Temporal Logic*, John Wiley & Sons, 1996.

Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, Vol. 37, No. 7, pp. 122-125.

Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2007). A Language and Environment for Analysis of Dynamics by SimulaTiOn. *Int. Journal of Artificial Intelligence Tools*, to appear, 2007. Short version in: Eymann, T. et al. (eds.), Proc. MATES'05. Springer LNAI, vol. 3550, 2005, pp. 165-178.

Burns, M., Bird, D., Leach, C., and Higgins, K. (2003). Anger management training: the effects of a structured programme on the self-reported anger experience of forensic inpatients with learning disability. *Journal of Psychiatric and Mental Health Nursing*, Vol. 10, pp. 569-577.

Damasio, A. (2000). *The Feeling of What Happens: Body, Emotion and the Making of Consciousness*. MIT Press.

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

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

Gross, J.J. (2001). Emotion Regulation in Adulthood: Timing is Everything. *Current directions in psychological science*, Vol. 10, No. 6, pp. 214-219.

Marinier, R.P., and Laird, J.E. (2004). Toward a Comprehensive Computational Model of Emotions and Feelings. In: *Proc. of the Sixth International Conference on Cognitive Modeling, ICCM'04*. Lawrence Erlbaum, Mahwah, NJ, pp. 172-177.

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.

Ochsner, K.N., and Gross, J.J. (2005). The cognitive control of emotion. *Trends in Cognitive Sciences*, vol. 9, pp. 242-249.

Port, R.F., and Gelder, T. van (eds.). (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press, Cambridge, Mass.

Thayer, J.F., and Lane, R.D. (2000). A model of neurovisceral integration in emotion regulation and dysregulation. *Journal of Affective Disorders*, Vol. 61, pp. 201-216.

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 (Serial No. 240), pp. 25-52.

Towl, G.J., and Crighton, D.A. (1996). *The Handbook of Psychology for Forensic Practitioners*. Routledge, New York.

Velasquez, J. (1997). Modeling Emotions and Other Motivations in Synthetic Agents. In: *Proceedings of the Fourteenth National Conference on Artificial Intelligence, AAAI'97*, pp. 10-15.