Matching Skin Conductance Data to a Cognitive Model of Reappraisal

Tibor Bosse[1], Jessica Brenninckmeyer[2], Raffael Kalisch[2], Christian Paret[2], and Matthijs Pontier[1]

[1]VU University Amsterdam, Department of Artificial Intelligence,
de Boelelaan 1081, Amsterdam, 1081HV, the Netherlands
tbosse@few.vu.nl, mpr210@few.vu.nl

[2]University Medical Center Hamburg-Eppendorf (UKE), Institute for Systems Neuroscience,
Martinistr. 52, 20249 Hamburg, Germany
Jessica.Brenninkmeyer@stud.uni-hamburg.de, rkalisch@uke.uni-hamburg.de, christian_paret@hotmail.de

Abstract

In the present paper we show that an existing mathematical model of emotion regulation can, if reduced to its reappraisal-specific components, fit skin conductance data obtained from an empirical study of reappraisal. By applying parameter tuning techniques, optimal fits of the model have been found against the (averaged) patterns of the skin conductance data. The errors that were found turned out to be relatively low. Moreover, they have been compared with the errors produced by a baseline variant of the model where the adaptive cycle has been removed, and were found substantially lower.

Keywords: emotion regulation, reappraisal, mathematical modeling, adaptation, skin conductance data.

Introduction

Emotion regulation refers to 'all of the conscious and nonconscious strategies we use to increase, maintain, or decrease one or more components of an emotional response' (Gross, 2001). This ability to regulate our own emotional states provides us with behavioral flexibility and is related to well-being and mental health (e.g., Gross, 1998, 2001; Ochsner and Gross, 2005; Thompson, 1994).

Recently, a number of authors have developed computational models of the processes related to emotion regulation and coping (e.g., Bach, 2008; Bosse et al., 2010; Gratch and Marsella, 2004; Marsella and Gratch, 2003; Reisenzein, 2009; Silverman, 2004). Computational models of emotion regulation may be useful for various reasons (see (Wehrle, 1998) for an overview). From a Cognitive Science perspective, they may provide more insights into the nature of affective disease and the working mechanisms of therapy. From an Artificial Intelligence perspective, they may be used to develop virtual agents with more human-like affective behavior.

In previous work (Bosse et al., 2010), we presented CoMERG, a Cognitive Model for Emotion Regulation based on Gross. Inspired by the theory put forward in (Gross, 2001), this model distinguishes five different strategies that humans typically use to affect their level of emotional response (for a given type of emotion) at different points in the process of emotion generation: *situation selection*, *situation modification*, *attentional deployment*, *cognitive change*, and *response modulation*. The different strategies and their effects are represented in the model via a set of difference equations.

An important asset of CoMERG is that the model is adaptive (see Bosse et al., 2007b). That is, based on the perceived success of an emotion regulation strategy that is performed, a person may adjust the degree of sensitivity of the process on the fly (e.g., in case a certain strategy does not decrease an undesired emotion sufficiently fast, the person may put more effort in the regulation). However, although a preliminary evaluation indicated that CoMERG produced plausible patterns (Bosse et al., 2010), to date the output of the model has never been compared with empirical data.

In order to assess to what extent CoMERG is able to reproduce empirical data, we here fit the model to skin conductance data that resulted from two empirical studies of reappraisal (unpublished material). Reappraisal, a variant of the cognitive change strategy aimed specifically at down-regulating emotion, is one of the most widely studied emotion regulation strategies. Gross (2001) defines reappraisal as a process where 'the individual reappraises or cognitively re-evaluates a potentially emotion-eliciting situation in terms that decrease its emotional impact'. For example, losing a tennis match is usually appraised as negative and would induce anger or sadness. To reduce these negative reactions, one could reappraise the situation by blaming the weather circumstances instead of the own capacities or by considering sportive success as irrelevant.

In (Kalisch, 2009), a novel (informal) model for reappraisal is presented, based on recent insights from imaging neuroscience. This model, called the implementation-maintenance model of reappraisal (IMMO), is characterized by its focus on the necessity of a mental reappraisal effort that needs to be maintained over the course of the emotional episode and is continuously adapted. Adaptation is realized through a loop of iterative evaluation and readjustment of the regulation process. IMMO thus shares a critical adaptation component with CoMERG.

To be able to better fit the results of CoMERG to the skin conductance data, the general model needs to be tailored specifically to reappraisal. Thus, the current paper has two main goals, namely 1) to refine the generic computational emotion regulation model CoMERG to the reappraisal

context, and 2) to evaluate the ability of the refined model to reproduce real data, by matching it to skin conductance data from empirical studies of reappraisal.

The remainder of this paper is structured as follows. First, the main mechanisms of CoMERG relevant to reappraisal are briefly summarized. Next, the setup of the reappraisal studies is described, with an emphasis on how the skin conductance data (to fit the new model) have been obtained. The following sections discuss how the model has been fit to the data, and present the results. The paper is concluded by a discussion.

CoMERG and its Extensions

CoMERG is composed of a set of difference equations, which represent how a person's emotional state changes based on certain regulation strategies. For convenience, the model concentrates on one specific type of emotion, in this case the fear induced by the threat of receiving a painful electric shock. We have chosen to express the *emotion response level* ERL in a real number, in the domain [0, 1]. A higher emotion response level means more fear.

In the model of Gross, five different elements n=1…5 (i.e., situation, sub-situation, aspect, meaning, and response) can influence the emotion response level. The experiments that produced the data to which the model is matched in this paper are restricted to the elements 1 (situation, i.e., threat of shock) and 4 (cognitive meaning, i.e., reappraisal). In the model, at any point in time, a certain *emotional value* v_n in the domain [0, 1] is attached to each element, representing the extent to which the current state of that element induces emotions. The model describes how persons increase or decrease those emotional values by comparing them with some desired values (or norms) v_{n_norm}. Because the participants receive explicit instructions about how to cognitively reappraise events, for element 4 we introduce an explicit $v_{4\text{-norm}}$ in the domain [0, 1]. A value of 0 for $v_{4\text{-norm}}$ would mean that one's aim is to reappraise the situation as not dangerous or frightening.

The emotional value contributes to the emotion response level ERL via an element-specific weight factor w_n, thereby taking into account a persistency factor β, indicating the degree of persistence or slowness of adjusting 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 seconds) will have a low β.

The regulation process of the cognitive meaning compares the actual cognitive meaning v_4 to v_{4_norm} at any time point. The difference d between the two is the basis for the adjustment of v_4. We assume that the self-monitoring process necessary to determine a deviation from $v_{4\text{-norm}}$ is a rather slow and effortful conscious process. We emulate this by the variable *eval* which is the integral of d over the past 3 seconds. Adjustment occurs via enhancing or reducing the cognitive effort made to achieve the desired emotional value $v_{4\text{-norm}}$, if *eval* signals a deviation. The regulation effort is expressed in the *modification factor* α_n (Bosse et al., 2007b), i.e., the 'willingness' to change the emotional value of element n. The effort one makes thus responds to a sort of reflection or meta-cognition about the emotion regulation process based on the history of differences d. One step further, the modification factor itself is adaptable as well: an additional *adaptation factor* γ_n represents the personal flexibility to adjust the emotion regulation behaviour (i.e., the personal tendency to adjust the emotional value of element n much or little). This depends on the cognitive costs of reappraising, which are represented by c_4.

The model is shown in a qualitative manner in the graph depicted in Figure 1. The variables above the dashed line represent the adaptation layer. The model without adaptation layer (Bosse et al., 2007a) will serve as a control condition to explore the necessity of this layer.

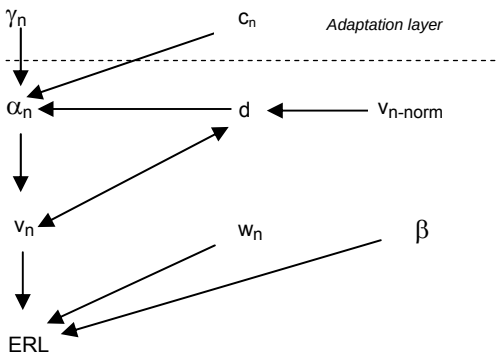

Figure 1: Dependencies between the Variables.

The main difference equations used to model these cycles are the following (see (Bosse et al., 2010) for more details):

Emotion Response Level

$$ERL(t+\Delta t) = (1 - \beta) * \Sigma_k(w_k * v_k(t)) + \beta * ERL(t)$$

Emotional Values

$$v_n(t+\Delta t) = v_n(t) - \alpha_n(t) * eval(t)/d_{max}$$

Modification Factors

$$\alpha_n(t+\Delta t) = \alpha_n(t) + \gamma_n * (\alpha_n(t) / (1 + \alpha_n(t))) * (abs\_eval(t) - c_n)$$

In terms of IMMO, determining eval can be seen as monitoring reappraisal success whose outcomes leads to an adjustment of the reappraisal effort α_4. Note the difference between eval (which is calculated by taking the integral of d) and abs_eval (which is calculated by taking the integral of the absolute value of d).

Obtaining the Data

To obtain skin conductance data about reappraisal processes, two within-subject experiments were performed. In both experiments, subjects were informed by an auditory warning signal that they might receive a shock to their hand at 25% probability during a given trial period (fear induction procedure). The warning signal was followed by another auditory cue telling them whether to reappraise (R)

the situation or not (NR). Generally, the reappraisal strategy consisted in taking a detached-observer perspective situation; that is, in distancing oneself from the situation and interpreting it as not affecting the core-self but being self-irrelevant. More specifically, in experiment 1 (n=24 right-handed healthy male subjects), subjects were told to imagine across both R and NR conditions a cloud in the sky that would symbolize the emotional aspects of a given situation, including all potential threats and accompanying reactions or feelings of tension or anxiety. For the R condition, they were asked to imagine themselves far away from this cloud, for example standing on a hill and observing the cloud from a distance (but not to look away). In addition to this mental image, they were given a self-statement that expressed the detached perspective: "The cloud is far out on the horizon. I observe it from a distance." For the NR condition, subjects were told to imagine themselves surrounded by the cloud and to use the corresponding self-statement: "I am inside the cloud. It surrounds me from all sides". They were to subvocally rehearse the appropriate statement throughout trials and to simultaneously keep the corresponding mental image in working memory. A similar strategy has been shown in previous studies to reduce fear of shock (Houston and Holmes, 1974; Kalisch et al., 2005). In Experiment 2 (n=20 right-handed healthy male subjects), the R condition was identical to experiment 1 whereas in NR trials, subjects were instead not told to use any self-statement or imagery but to simply attend to the situation and allow their emotional reaction to unfold, not attempting to change it.

Skin conductance is a measure of the sympathetic arousal that accompanies most fear responses. Although it cannot capture all aspects of a fear response, it is one of the few available continuous and objective measures of the response and was thus used to generate ERL time courses.

In all figures below, skin conductance time courses are averaged across trials and subjects in that experiment. Solid red lines represent average NR time courses, dotted red lines represent average R time courses.

Matching Data to the Model

To obtain a close fit of the simulation model to the empirical data obtained in the experiments, parameter tuning was used (Sorenson, 1980). Since the challenge is to tune the parameters of an existing dynamic model, rather than to come up with an optimal function from scratch, it is not possible to apply standard regression techniques in this case. Therefore, a dedicated parameter estimation method was used, which is similar to the approach used in (Bosse, Memon, Treur, and Umair, 2009). According to this approach, to match the model to the data it is first needed to obtain the sensitivity of a parameter: the change in difference between the model and the data with a given change in parameter value.

To determine the *sensitivity S*, a small change ΔP in the parameter is tried to make an additional prediction for X, and based on the resulting change ΔX found in the two predicted values for X, the sensitivity S can be estimated:

$$S_{X,P} = \Delta X / \Delta P$$

After the sensitivity is determined, a better guess for the value of P can be determined by taking

$$\Delta P = -\lambda * \Delta X / S_{X,P}$$

where ΔX is the deviation found between observed and predicted value of X; so, for example, when $\Delta X = 0.25$ and $\lambda = 0.3$, then for $S_{X,P} = 0.75$ this obtains $\Delta P = -0.3*0.25 / 0.75 = -0.1$. However, when the sensitivity $S_{X,P}$ is a bit smaller, it could be possible that the adjustment of the value of P based on the formula above would exceed the maximum or minimum value of its range. If this happened, the parameter was adjusted by intuition.

Based on this adjustment approach, the overall parameter tuning process was done as follows:

1. Take G the set of parameters P for which adjustment is desired; the other parameters are kept constant.
2. Assume initial values for all parameters P, and for λ.
3. By simulation determine predicted value CV_X at time point t for X, using the assumed values of the parameters.
4. For each parameter P in G, by simulation determine predicted value for V_X at time point t, using only for P a value changed by some chosen ΔP and the unchanged assumed values for the other parameters.
5. For each parameter P in G determine the sensitivity $S_{X,P}$ of X for P at time point t by dividing the difference between values for X found in step 4 and 5 by ΔP:
   $$S_{X,P} = (CV_X - V_X) / \Delta P$$
6. For each parameter P determine the change ΔP as
   $$-\lambda * \Delta X / S_{X,P}$$
7. For each parameter P adjust its value by ΔP.
8. Return to step 1 until the fit is satisfactory.

The coefficient of determination R^2 (Steel & Torrie, 1960) was calculated to determine the quality of the fit (the closer to 1 the better). The match was called satisfactory when the quality of fit did not increase anymore for several time steps. If the matching process seemed to be stuck into a local optimum, the parameters were adjusted by intuition to check whether the match could be improved.

The set of parameters G looked at were β, γ, c, α, and w_1. We did not use any constraints for the values, except that w_1 should always be bigger than w_4, as Gross described that emotion regulation strategies performed earlier in the regulation process are more effective (Gross, 2001).

Results

In this section, the results of the skin experiments are described, as well as the curves produced by fitting the model on the results. For both experiments, first the fits produced by the complete model (with adaptation) are

presented, both for the NR and for the R condition, followed by the fits produced by the model without adaptation (with was used as a control condition).

Exp1 – Adaptation – No Reappraisal (NR)

We modeled the NR condition (solid line in the figures) by setting v_{4_norm} to the same level as v_1 and v_4 (which is always = v_1 at the start of the simulation). This models that subjects do not intend to change their appraisal of the situation but allow their automatic appraisal systems to dominate and thus to solely determine the ERL.

Because v_{4_norm} has the same value as v_4, d = 0, and v_4 is not changed during the experiment. Therefore, α_4 has no influence on v_4, and thus no indirect influence on ERL. For the same reason, c_4 and γ_4 have no indirect influence on the ERL. Further, since v_1 and v_4 have the same value throughout the complete experiment, the proportion of w_1 does not influence ERL either. This leaves the parameter β_{ERL} as the only possible factor for fitting the data.

Using the method described earlier in this paper, the optimal fit to the data was found for $\beta_{ERL} = 0.9841$. The R^2 of the fit was 0.9960, and can be seen in the higher curve of Figure 2.

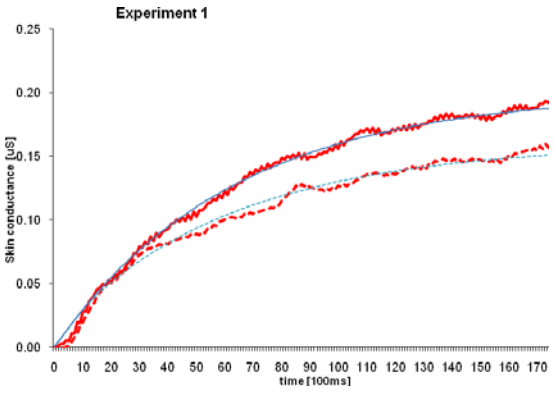

Figure 2: The fits of the model with the adaptation layer to Experiment 1. Thick, red, solid line: average skin conductance from Non-Reappraisal (NR) trials; Thick, red, dotted line: Reappraisal (R) trials. Thin, blue, solid line: model fit for NR trials; Thin, blue, dotted line: model fit for R trials.

Exp1 – Adaptation – Reappraisal (R)

The goal to reappraise the situation as not self-relevant was modeled by setting $v_{4\_norm} = 0$. The starting value of v_4 was still modeled to be the same as v_1. Because this creates a discrepancy d between v_4 and v_{4_norm}, now all the five parameters have a direct or indirect influence on ERL. However, because β_{ERL} represents a personality factor which shouldn't differ among experimental conditions, the value from the NR fit above was taken. This leaves the other four parameters for optimizing the fit.

Using the method described earlier in this paper, the optimal fit to the data was found for the following parameter settings:

$\alpha_4 = 0.188$
$w_1 = 0.6$
$\gamma = 0.2$
$c = 0.4$

This led to a fit with $R^2 = 0.9876$, which can be seen in the lower curve of Figure 2.

Exp1 – No Adaptation – Reappraisal (R)

To explore the necessity of the adaptation layer in the emotion regulation model, we also made a fit for the model without the adaptation layer, in which α is kept constant. Because the fit for the NR condition already had a constant α, the curve does not change.

Because γ and c are part of the adaptation layer, they cannot be considered for fitting the R condition, leaving α and w_1 for optimizing the fit.

The optimal fit to the R data was found for the following parameter settings:

$\alpha_4 = 0.027$
$w_1 = 0.79$

As can be seen in the lower curve of Figure 3, the fit still is still reasonable, with an R^2 of 0.9723. However, it was worse than the fit of the model with the adaptation layer added, where an R^2 of 0.9876 was reached.

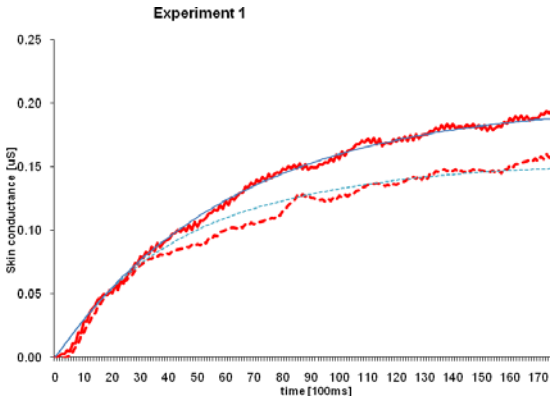

Figure 3: The fits of the model without the adaptation layer to Experiment 1.

Exp2 – Adaptation – No Reappraisal (NR)

In experiment 2 in NR trials, participants were instructed to think or feel as they normally would in such a situation. No

cognitive effort to maintain any type of statement or image was required. This was modeled by setting $\alpha_4 = 0$.

Because the update mechanism of α_4 is proportional to α_4, it would always stay at 0. Therefore, γ, and c had no direct or indirect influence on ERL, and were not considered. Because α stayed at 0 throughout the experiment, v_4 also stayed constant, at the same level as v_1. Therefore, w_1 also did not influence ERL, leaving only β_{ERL} for optimizing the fit to the data.

The optimal fit, which can be seen in the higher curve of Figure 4, was found for $\beta_{ERL} = 0.9869$, with an R^2 of 0.9556.

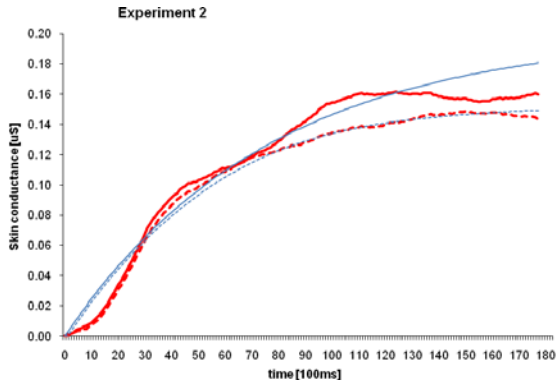

Figure 4: The fit of the model with the adaptation layer to Experiment 2.

Exp2 – Adaptation – Reappraisal (R)

In the R condition, the value for β_{ERL} was taken from the value found in the NR condition, and the other four parameters could all be used for optimizing the fit to the data, similar to the R condition of experiment 1.

The optimal fit was found for the following parameter settings:

$\alpha_4 = 0.003$
$w_1 = 0.75$
$\gamma = 0.3$
$c = 0.1$

This led to a fit with $R^2 = 0.9818$, which can be seen in the lower curve of Figure 4.

Exp2 – No Adaptation – Reappraisal (R)

For experiment 2 we also made a fit with the model without the adaptation layer. Again, because γ and c are part of the adaptation layer they cannot be considered for making the fit, leaving α_4 and w_1 for optimizing the fit.

The optimal fit to the data was found for the following parameter settings:

$\alpha_4 = 0.004$
$w_1 = 0.51$

As can be seen in Figure 5, the fit is still quite good, with an R^2 of 0.9806, but slightly worse than could be made using the version of the model with the adaptation layer, with which an R^2 of 0.9818 was reached.

These results illustrate that the emotion regulation model by (Bosse et al., 2010) is capable of reproducing empirical data quite closely. Moreover, the fact that the fits of the model without the adaptation layer are worse provide evidence that reappraisal as performed by humans may indeed be an adaptive process.

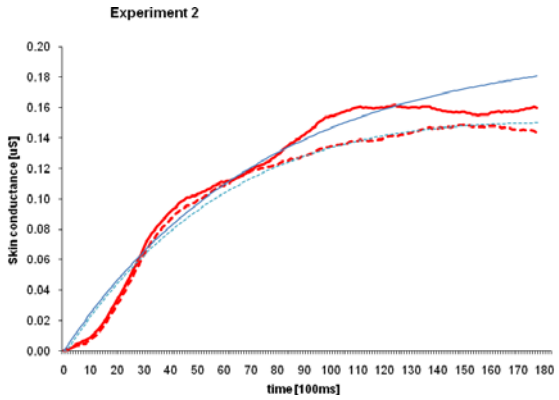

Figure 5: The fit of the model without the adaptation layer to Experiment 2.

Discussion

Over the last decade, the number of computational models of affect has rapidly increased, especially in the area of Artificial Intelligence (e.g., Bach, 2008; Bosse et al., 2010; Gratch and Marsella, 2004; Marsella and Gratch, 2003; Reisenzein, 2009; Silverman, 2004). Most of these models have as their main goal to endow virtual agents (e.g., robots or avatars) with more believable human-like behavior. However, only a small subset of these approaches aims to reproduce the dynamics of the more subtle sub-processes involved (such as reappraisal) in a detailed manner (see (Bosse et al., 2010), for an extensive literature overview). An even smaller subset validates the results of the model against physiological data, such as skin conductance or fMRI data, yielding a large gap between AI-inspired modeling approaches and empirical psychological research.

The main contribution of the present paper is a first step towards closing this gap. We have shown that an existing cognitive model of emotion regulation can, if reduced to its reappraisal-specific components, fit empirical data. By applying parameter tuning techniques, optimal fits of the model have been found against the (averaged) patterns of the skin conductance data. The errors that were found turned out to be relatively low. Moreover, they have been compared with the errors produced by a baseline variant of the model where the adaptive cycle has been removed, and

were found substantially lower. Although this is obviously not an exhaustive proof for the correctness of the model, it is an important indication that reappraisal as performed by humans may indeed be an adaptive process, as has been postulated by current informal models of reappraisal (Kalisch, 2009).

Further validation and refinements of our model are obviously warranted. Regarding validation, the current work should be seen as an initial test whether the CoMERG model is capable of reproducing empirical data at all. In future research, more extensive tests will be performed, based on cross-validation and involving more participants. Regarding model refinement, it will be particularly interesting to see whether it can be adjusted to also simulate a proposed subparcellation of reappraisal effort into an early retrieval and a later working memory maintenance and monitoring component that has ensued from a recent analysis of neuroimaging data (Kalisch, 2009). The model might then also be useful for prediction brain activation time courses.

Acknowledgments

We would like to thank Jan Treur for a number of fruitful discussions.

References

Bach, J. (2008). Principles of Synthetic Intelligence: Building Blocks for an Architecture of Motivated Cognition. New York, NY: Oxford University Press.

Bosse, T., Memon, Z.A., Treur, J., and Umair, M. (2009). An Adaptive Human-Aware Software Agent Supporting Attention-Demanding Tasks. In: Yang, J.-J., Yokoo, M., Ito, T., Jin, Z., and Scerri, P. (eds.), *Proceedings of the 12th International Conference on Principles of Practice in Multi-Agent Systems, PRIMA'09*. Lecture Notes in AI, vol. 5925. Springer Verlag, 2009, pp. 292-307.

Bosse, T., Pontier, M., and Treur, J. (2007a). A Dynamical System Modelling Approach to Gross´ Model of Emotion Regulation. In Lewis, R.L., Polk, T.A., and Laird, J.E. (eds.), *Proceedings of the 8th International Conference on Cognitive Modeling*, ICCM'07. Taylor and Francis, 2007, pp. 187-192.

Bosse, T., Pontier, M., and Treur, J. (2007b). A Computational Model for Adaptive Emotion Regulation. In Lin, T.Y., Bradshaw, J.M., Klusch, M., Zhang, C., Broder, A., and Ho, H. (Eds.) *Proceedings of the Sixth IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT'07*. IEEE Computer Society Press, 2007, pp. 289-293.

Bosse, T., Pontier, M., and Treur, J. (2010). A Computational Model based on Gross' Emotion Regulation Theory. *Cognitive Systems Research Journal*, vol. 11, 2010, pp. 211-230.

Gratch, J., and Marsella, S. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research*, vol. 5, issue 4, 2004, pp. 269-306.

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.

Houston, B., Holmes, D.S. (1974). Effect of avoidant thinking and reappraisal for coping with threat involving temporal uncertainty. *J. Pers. Soc. Psychol.*, 30, pp. 382-388.

Kalisch, R. (2009). The functional neuroanatomy of reappraisal: Time matters. *Neuroscience and Biobehavioral Reviews* 33, pp. 1215-1226.

Kalisch, R., Wiech, K., Critchley, H.D., Seymour, B., O'Doherty, J.P., et al. (2005). Anxiety reduction through detachment: subjective, physiological, and neural effects. *J. Cogn. Neurosci.*, 17, pp. 874-883. Kosslyn, S.M., Thompson.

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.

Reisenzein, R. (2009) Emotions as metarepresentational states of mind: Naturalizing the belief-desire theory of emotion. *Cognitive Systems Research*, vol. 10, pp. 6-20.

Silverman, B.G. (2004). Toward Realism in Human Performance Simulation. In: Ness, J.W., Ritzer, D.R., and Tepe, V. (eds.), The science and simulation of human performance, pp. 469-498. Amsterdam, Elsevier.

Sorenson, H.W. (1980). Parameter estimation: principles and problems. Marcel Dekker, Inc., New York.

Steel, R.G.D. and Torrie, J.H. (1960), Principles and Procedures of Statistics, New York: McGraw-Hill, pp. 187-287

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.

Wehrle, T. (1998). Motivations behind modelling emotional agents: Whose emotion does your robot have? In C. Numaoka, L. D. Cañamero & P. Petta (Eds.), Grounding Emotions in Adaptive Systems. Zurich: *5th International Conference of the Society for Adaptive Behavior Workshop Notes (SAB'98).*