

SuperDreamCity: An Immersive Virtual Reality Experience that Responds to Electrodermal Activity

Doron Friedman¹, Kana Suji², and Mel Slater³

¹ Interdisciplinary Center, Herzliya, Israel,
doronf@idc.ac.il,

² Dream Products Co.,
dreamproductsco@yahoo.co.uk

³ ICREA-Universitat Politecnica de Catalunya, Spain
and Department of Computer Science, UCL
m.slater@cs.ucl.ac.uk

Abstract. In this paper we describe an artistic exhibition that took place in our highly-immersive virtual-reality laboratory. We have allowed visitors to explore a virtual landscape based on the content of night dreams, where the navigation inside the landscape was based on an online feedback from their electrodermal response. We analyze a subset of the physiology data captured from participants and describe a new method for analyzing dynamic physiological experiences based on hidden Markov models.

1 Introduction

This study is part of a research that assumes an experimental paradigm where a person is exposed to stimuli that induce physiological changes (such as changes in heart rate (HR), heart rate variability (HRV), electrodermal activity (EDA), and similar autonomous responses). A computer program monitors how the physiology changes over time and in response to sequences of visual stimuli. The automated decisions related to the presentation of the visual stimuli are planned to have some desired impact on the participant's physiological state.

Such research could be considered complementary to traditional biofeedback. "Classic" biofeedback involves measuring a subject's bodily processes such as blood pressure or galvanic skin response (GSR) and using a machine to convey this information to him or her in real-time in order to allow him or her to gain control over physical processes previously considered automatic [3, 9]. Biofeedback thus has a number of therapeutic uses in helping people learn how to achieve and control positive mental states such as concentration or relaxation, and has been used with people with anxiety, depression and attention problems [19]. Our view is that we can now revisit traditional biofeedback taking into account advances in online signal processing, intelligent computation, and various types of

feedback, such as, in this case, highly-immersive virtual reality (VR). Our approach is almost the inverse: in our case the machine is the one supposed to do the learning and adaptation, and not the person.

In this study we report on an early step where we integrated a highly immersive Cave-based experience with real-time feedback based on skin conductance [1]. While this was not a scientifically controlled experiment, we show how the results can be systematically analyzed.

2 Background

GSR, also sometimes called electrodermal activity (EDA), is measured by passing a small current through a pair of electrodes placed on the surface of the skin and measuring the conductivity level. Skin conductance is considered to be a function of the sweat gland activity and the skin's pore size. The real-time variation in conductance, which is the inverse of the resistance, is calculated. As a person becomes more or less stressed, the skin's conductance increases or decreases proportionally [1]. There are two measures associated with GSR: one is overall level, called the tonal level, which gives the overall level of arousal, and the other is skin conductance response (SCR), which gives arousal in response to specific events (or unknown random internal events). In our study we have used the tonal level.

The idea of closed-loop VR has already been addressed by the sci-art community. One of the classic VR art pieces of all times is *Osmose* [5], where the participants' experience depends on the analysis of their breathing. Another, more recent art piece related with body-centered interaction in VR include *Traces* by Simon Penny⁴. These art projects are highly influential in raising discussions regarding interface design practices. However, there is no attempt for any scientific analysis of the experience, in terms of the human-machine feedback loop, and no analysis of the data. Some interactive applications or games using biofeedback have proved useful for relaxation (as an example based on EEG see [10]).

We have come upon such man-machine loop issues in our recent studies in brain-computer interfaces (BCI) in highly-immersive VR [8, 14]. Such BCI includes training human subjects to control a computer system by "thought", based on real-time analysis of electroencephalogram (EEG). It involves two complex, interdependent systems: the brain and the machine, and in order for the BCI to be successful they both need to learn. The solution typically adapted, is to allow each of the systems to learn in separate, while the other is kept constant [13]. The research proposed here similarly suggests studying this issue of mutual adaptation, but in a different context.

Picard [15] coined the term *affective computing*: this includes computers that both recognize and exhibit emotions. Picard, as well others in this area of research, have demonstrated devices based on real-time analysis of autonomic responses, such as: affective jewelry and accessories, affective toys, affective

⁴ <http://www.medienkunstnetz.de/works/traces>

tutoring systems, computer responses to user frustration, and visualization of the user’s emotional state [15]. Recognition of emotions is addressed by several means, physiological responses being one of them.

Bersak et al. coined the term *affective feedback*, which means that “the computer is an active intelligent participant in the biofeedback loop” [2]; where both player and game are affected by the actions of the other. Prendinger and his colleagues have developed and evaluated a closed-loop virtual agent that responds to users’ emotions. The valence and intensity of emotions are recognized based on skin-conductance level and electromiography [17, 18, 16].

The so-called *affective loop* has also been described by Hook and colleagues; see for example [22]. It has been shown in systems like SenToy [12], eMoto [22], Affective Diary [11] and Brainboll [21] that it is indeed possible to involve users in affective loops, but that the design needs to be carefully crafted to the specific demands of the application functionality in order for the application to work.

3 The VR Experience as an Experiment

3.1 Scientific Objective

The objective of the study is to test whether the physiological state of a VR participant may be manipulated systematically over time, during a VR experience. In addition, we suggest methods for analyzing the data and inspecting whether the manipulation was achieved.

Such intelligent systems for physiological manipulation may be based on several computation paradigms.

Our approach in this paper is based on reinforcement loops – Such an approach would try to use positive and negative feedback loops; these were investigated as early as the middle of the twentieth century [23]. Positive loops may be used to drive an existing trend to an extreme, and negative loops may be used to extinguish existing trends.

Specifically, our assumption is that we can induce positive feedback loops by leading participants into positive spaces when they are relaxed and into negative spaces when they are stressed (or aroused). If the system is successful, we would see two types of patterns: in one case participants will mostly visit positive spaces, and their overall GSR levels would remain flat, or even decrease. In the other case, participants would mostly visit negative places and their overall GSR level will increase significantly during the experience.

This assumption can be broken into two hypotheses:

1. Negative places would have a significantly different impact on GSR tonal level than positive places – specifically, the GSR level would increase after negative places and decrease after positive places; and
2. An analysis of the dynamics of transitions between positive and negative places would reveal the existence of positive feedback loops.

3.2 The VR System

The study was carried out in a four-sided ReaCTor system that is similar to a Cave [4]. The UCL Cave is a 2.8x3x3 meter room with stereo projection on three walls and on the floor. The participant wears light-weight shutter glasses and an Intersense IS900 wireless head-tracker. The result is that the participant is free to move around the room and is (almost) surrounded by the virtual landscape.

3.3 The Virtual Environment

The content of the virtual environment (VE) is based on work by the second author, who is a London-based artist. She is in the (fictional) business of buying dreams: she pays people one Great British Pound each so that they tell her about their night dreams. Then she models the dreams in 3D, and adds them into DreamCity – an online version, where people are able to browse among other people’s dreamscapes (<http://www.dreamproductsco.com>).

For the London Node (Networked, Open, Distributed Event) media-art festival, March 2006, we decided to create a unique version of DreamCity, called SuperDreamCity. First, rather than displaying the models on a desktop computer, we adapted DreamCity for the Cave. Second, we decided that the participants will explore the dreamscape using their physiological responses.

For SuperDreamCity the second author selected several “dreams” into one VE where all the dreamscapes were randomly scattered around (see Figure 1); we have only used static models in this version. Most of the dream sites includes sound files that played when the participant was in the site vicinity. The VE included a low-volume background music playing in a loop, for the purpose of “atmosphere building” – this was a dream-like electronic music (by musician Laurie Anderson).⁵

3.4 Real-time Physiology

We wanted to allow the participants to explore the VE in a way that would depend on their internal bodily responses to the environment, as reflected in their autonomous nervous-system responses. We have selected GSR as a single measurement, since this is easily measured by a small sensor placed on two fingers, which is easy and quick to fit; this was important as we were attempting a quick turnover of visitors. We have used the raw GSR values (the tonal GSR level) as a single feature in affecting the navigation.

We have carried out previous work in real-time neurophysiology in the Cave [8]. It was relatively straightforward to convert the system to use for real-time GSR. In this case we used the g.Mobilab system (g.Tec, Austria), which includes sensors, a small amplifier, and software. GSR was sampled at 32 Hz, and the signal was obtained from electrodes on two fingers. The g.Mobilab software is easy to modify – it includes a Matlab/Simulink model for the device. We have

⁵ A video is available online in <http://www.cs.ucl.ac.uk/staff/d.friedman/sdc/sdc.mov>



(a)



(b)

Fig. 1. (a) A screenshot of an industrial area from a dream, as viewed online. (b) A participant in the VR Cave experiencing the same industrial area in SuperDreamCity. Note that this image is for illustration: in the actual experience the participant would not be holding the navigation wand (as they navigate based on GSR) and the image would be stereoscopic.

used Simulink to extract the raw GSR value, pass it to a dynamically-linked library (DLL) and over the network using the Virtual Reality Peripheral Network (VRPN)⁶. On the Cave Irix system, a VRPN client would intercept the raw GSR values and feed them into the VR application. The VR software was implemented on top of the DIVE software [7, 20]. The DIVE application would then implement the navigation logic based on the real-time GSR value (this is scripted in TCL).

3.5 Method

For the show, the artist re-created the VE with 20 of her own dreams modeled in 3D, 10 having positive associations and 10 having negative associations. As an example of a positive dream consider an amusement park, and as a negative dream consider industrial areas. The emotions were expressed with choice of colors and sound effects. In this case the emotional interpretation of the dreams was given by the artist or by the dreamer; clearly, in a controlled scientific experiment, this emotional interpretation needs to be validated.

Rather than a low-level mapping of GSR into navigation, we have opted for a high-level mapping. We decided to split the experience into stages. First, all subjects find themselves floating over one of the positive dream sites. Then, in each stage of the experience they start floating from one dream site towards another site. The decision to what site to navigate is based on the trend of the GSR⁷.

For the art exhibition, we decided to explore positive feedback loops, i.e., the system would try to reinforce the participant's physiological trend. If the GSR value increased from the previous section, a random negative dream site was targeted. If overall GSR decreased, a random positive site was selected. Navigation speed was also modified — for every selection of a negative site the speed was increased by 10% of the baseline speed, and, correspondingly, for every selection of a positive site the speed decreased by 10%. Thus, our expectation was that this VE would create a positive feedback loop with the participant — i.e., we expected that some participants will keep visiting negative sites, which would increase their GSR, so that overall they would mostly visit negative sites and become increasingly stressed throughout the experience. We expected that for other subjects there would be a relaxation loop, such that their GSR would gradually decrease as they keep visiting positive sites and floating in a slower and more relaxed fashion. In the next session we explore how this was evaluated scientifically, and report the results.

⁶ <http://www.cs.unc.edu/Research/vrpn/>

⁷ In states of increased excitement people sweat more, which should result in a higher GSR as compared with a relaxed state.

4 Experimental Procedure

Our assumption was that, under some conditions, an exhibition open to the public can serve as a scientific experiment (for another example see [6]); in the least case, the data collected can serve as useful insight for future research.

The London Node Festival took place over a whole month, included dozens of events in different locations around the city, and advertised online. We have advertised our exhibition, in our VR lab, to be open to the public for a few hours each day over three consecutive days (over the weekend), and required people to register in advance online. Each registered person received a time slot to show up in the lab (with 20 minutes allocated per person).

During the exhibition there were at least three people working in the lab. One person was necessary to fit the GSR device and operate the systems. Another person stayed outside the lab space, and managed the queue of people. Finally, the artist greeted each person into the experiment. She was dressed as a businesswoman, handed them her business card, explained to them about her (fictional) business buying dreams, and explained to them that they are about to experience a dreamscape that would respond to their physiology.

When participants were led into the Cave room they were fitted with the GSR sensor and goggles, and placed inside the dark Cave. They were instructed to wait there. Then there was a period of at least 60 seconds, after which the VR experience began – this duration was used for measuring GSR baseline. Participants stayed in the Cave for varying durations of 5-15 minutes, based on the queue outside. Most participants loved the experience and would have stayed more if they were allowed.

5 Results

During the three exhibition days we had 35 participants in the Cave. We collected data for all participants, but most of the sessions had to be discarded. Because this was an art exhibition, participants behaved in quite different ways than subjects would behave in a typical scientific experiment in our Cave. Some of them talked a lot, moved a lot, tried to jump, or even, in one case, lie down on the Cave floor. In some cases we had a long queue outside and had to allow more than one person into the Cave. All these sessions were discarded. Out of the remaining sessions, 15 sessions included valid GSR data (these were most of the “good” sessions), and these were analyzed as described below.

Each session is characterized by a number of events – an event is the point in time when the system decided to navigate into another dream site, either negative or positive. The duration elapsing between two events varies, as it depends on varying navigation speeds and on variable distances among the pairs of dream sites. The duration between events was always at least 20 seconds, sometimes up to one minute. Thus each session included a different number of events, ranging from 7 to 35.

First, we want to test whether positive dreamscapes affect GSR in a different way than negative dreamscapes and examine the trends in GSR tonal level

around the events. This is tested using an analysis of covariance (ANOCOVA). We take the time around the events (from 20 seconds before the event to 20 seconds after the event) to be the predictor x , the GSR level to be the response variable y , and a binary variable c for the dream category. If our hypothesis is correct then we expect the coefficient of the positive dreams to be significant with a negative slope, the coefficient of the negative dreams to be significant with a positive slope, and the Anova value for $x \cdot c$ to be significant.

A case by case study reveals that the hypothesis was correct for 5 out of the 15 subjects: cases where the slope was significantly different between the two events, and the trend for negative dreams was higher than for positive dreams (this includes cases such as in Figure 2, where both trends were decreasing, but the positive dreams decreased faster). For 9 subjects the results were not significant, and for one subject the results were significant, but they were the opposite of our prediction: the positive dreams resulted in an increase in GSR and the negative dreams in a decrease.

After normalizing the GSR values for all subjects, we can perform the same analysis for the data taken from all subjects together. Our hypothesis is not supported, i.e., the experience, taken over all subjects, did not cause increase and decrease in GSR levels as predicted.

Our main interest is in the dynamics of the experience. Since the first hypothesis was not fully supported we did not expect to find the dynamics we expected, but we still describe how we suggest to analyze such data. We model each session as a stochastic process over state transitions. There are two states: P (positive) and N (negative), according to the two types of dreams. Accordingly, there are four types of transition types: PP, PN, NP, and NN. Furthermore, we can distinguish between two types of transitions: T transitions that keep the current trend (NN and PP) and R transitions – trend reversal transitions (NP and PN).

Figure 3 illustrates that, indeed, the state transitions seem random. More formally, the data from each session can be modeled as a hidden Markov model (HMM): we observe a sequence of emissions, and our goal is to recover the state information from the observed data.

Our HMM includes two states: P and N. We know the emission matrix for the model: when the system is in state P there is a probability of 0.1 for events 1 – 10 to occur and a probability of 0 for events 11 – 20 to occur. Conversely, when the system is state N there is a probability of 0 for events 1 – 10 to occur and a probability of 0.1 for events 11 – 20 to occur. For each session we know the state path and the emission sequence. Based on these parameters we can estimate the transition matrix, which is the only unknown parameter, for each session. For each transition from state S_1 to state S_2 the estimation of the transition probability is given by the number of transitions from S_1 to S_2 in the sequence divided by the total number of transitions from S_1 in the sequence.

In our case there are two states only, so the transition matrix has two free parameters: if we denote by α the probability for a transition from P to P then clearly the probability for moving from P to N is $1 - \alpha$; similarly, we denote by

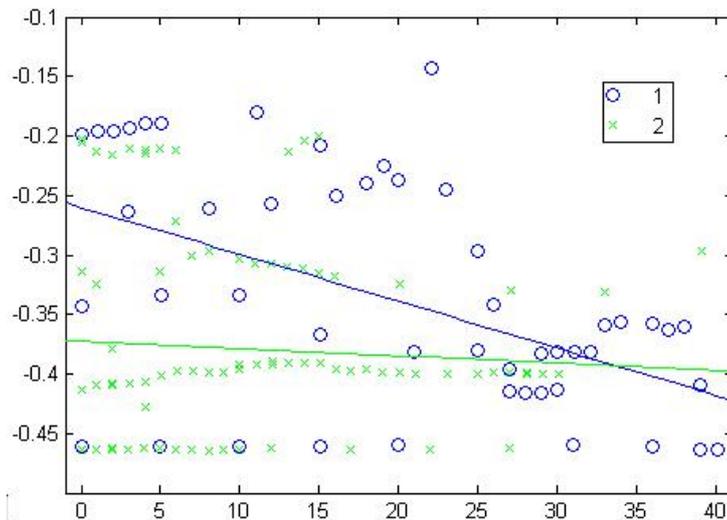


Fig. 2. The analysis of variance plot for one subject, showing GSR as a function of time around the events. In this case we see that both event categories resulted in a decreasing trend of GSR, but positive (blue) decreased more than negative (green). For this subject the difference is significant. Note that we do not care about the intercept of the regression line, only the slope (since each event starts in a different level). To make the results apprehensible each is the average of 50 GSR samples.

β the probability for a transition from N to P and then the probability for a transition from N to N is $1 - \beta$.

Thus, For each session we have two observations, resulting in two response variables: α and β . In our case α is in the range $0 - 0.8$ with mean 0.38 and β is in the range $0.25 - 1$ with mean 0.63 ⁸. Most importantly, an Anova test reveals that for both variables we cannot reject the null hypothesis, i.e., we have an indication that both α and β are random. We note that trying to estimate the emission matrix of our model does result in a rejection of the null hypothesis, i.e., the probabilities for selecting an event based on a state are not arbitrary. This indicates that our analysis should have revealed a pattern in the transition matrix if there was one.

⁸ The fact that $\alpha + \beta \cong 1$ is only a coincidence.

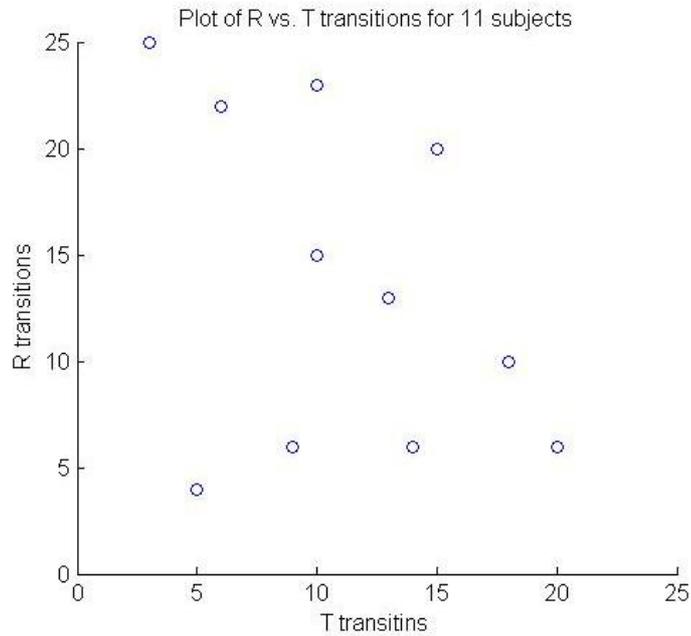


Fig. 3. A plot of the ratio of the state-preserving transitions (PP and NN) vs. the state-change transitions (PN and NP) for 11 out of 15 subjects (for four subjects the number of transitions in the session was too small). We see that the scatter seems uniform inside the lower left triangle part of the space, as expected from a random process. If, as we expected, the experience would have had a positive-loop impact, we would expect the points to be concentrated near the axes of the diagrams. If the experience would have enforced negative feedback loops we would have expected the points to be around the diagonal $y = x$.

6 Discussion

We are interested in studying the dynamics of human physiology when participants are placed inside immersive environments that respond to this physiology. We have described how this dynamics was implemented and studied in the scope of an artistic exhibition.

There is growing interest in such affective-loop systems, for various applications, including training, psychological treatment and entertainment. However, the dynamics of such closed-loop systems is rarely studied in a systematic way. It would be of both theoretic and practical interest to have a better understanding of the way media systems, providing sensory inputs, may affect people's autonomous responses over time, especially in the context of a closed-loop system.

One of our main lessons from this study is that while it is now feasible to create this type of biofeedback application, even using highly-immersive VR, it is not easy to create a meaningful experience that fully exploits the possibilities of biofeedback in highly-immersive VR. In our case our analysis revealed that the feedback loop did not take place as expected (unless, possibly, for 5 out of 15 subjects whose data was analyzed). For most participants the biofeedback part of the experience was probably meaningless, in the sense that the experience had no systematic effect on the participant's physiology. This is probably the case in many similar art projects, but these do not even report the results, let alone analyze the data.

There are several lessons and ways to go forward. For example, raw GSR is not necessarily the best feature to use for such neurophysiological experiences. It is probably better to use SCR (the number and/or amplitude of peaks in GSR as a response to a new stimulus), heart rate, or some combination of these features.

As a result of this study, we are currently revisiting the same questions, using a similar approach, in the context of a more scientific methodology. Obviously, such experiments would first validate the effects of the selected stimuli, before studying their dynamics. We suggest studying such experiences, based on real-time physiology, and analyze the degree of success using HMMs.

ACKNOWLEDGEMENTS

This work has been supported by the European Union FET project PRESENCIA, IST-2006-27731. We would like to thank David Swapp and other members of the VECG lab in UCL for their support. We would also like to thank Christoph Guger for his support with the gMobilab system.

References

1. J. L. Andreassi. *Psychophysiology: Human Behavior & Physiological Response*. Laurence Elbaum Associates, 2000.
2. D. Bersak, G. McDarby, N. Augenblick, P. McDarby, D. McDonnell, B. McDonal, and R. Karkun. Biofeedback using an immersive competitive environment. In *Online Proceedings for the Designing Ubiquitous Computing Games Workshop, Ubicomp 2001*, 2001.
3. E. B. Blanchard and L. D. Young. Clinical applications of biofeedback training: A review of evidence. *Archives of General Psychiatry*, 30:573–589, 1974.
4. C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: Audio visual experience automatic virtual environment. *Comm. ACM*, 35(6):65–72, June 1992.
5. C. Davies and J. Harrison. Osmose: Towards broadening the aesthetics of virtual reality. *ACM Computer Graphics [special issue on Virtual Reality]*, 30(4):25–28, 1996.
6. K. Eng, D. Klein, A. Babler, U. Bernardet, M. Blanchard, M. Costa, T. Delbruck, R. J. Douglas, K. Hepp, J. Manzolli, M. Mintz, F. Roth, U. Rutishauser, K. Wassermann, A. M. Whatley, A. Wittmann, R. Wyss, and P. F. M. J. Verschure. Design

- for a brain revisited: The neuromorphic design and functionality of the interactive space Ada. *Reviews in the Neurosciences*, 14:145–180, 2003.
7. E. Frecon, G. Smith, A. Steed, M. Stenius, and O. Stahl. An overview of the COVEN platform. *Presence: Teleoperators and Virtual Environments*, 10(1):109–127, Feb. 2001.
 8. D. Friedman, R. Leeb, C. Guger, A. Steed, G. Pfurtscheller, and M. Slater. Navigating virtual reality by thought: What is it like? *Presence: Teleoperators and Virtual Environments*, 16(1):100–110, 2007.
 9. K. R. Gaarder and P. Montgomery. Scientific foundation of biofeedback therapy. In K. R. Gaarder and P. Montgomery, editors, *Clinical Biofeedback*, pages 3–30. Williams & Wilkins, 1981.
 10. S. I. Hjelm, E. Eriksson, and C. Browall. BRAINBALL - Using brain activity for cool competition. In *Proc. First Nordic Conf. on Human-Computer Interaction 2000*, 2000.
 11. M. Lindstrom, A. Stahl, K. Hook, P. Sundstrom, J. Laakso, M. Combetto, A. Taylor, and R. Bresin. Affective diary - Designing for bodily expressiveness and self-reflection. In *Proc. ACM SIGCHI Conf. Computer-Human Interaction*, 2006. Work in Progress paper.
 12. A. Paiva, R. Chaves, M. Piedade, A. Bullock, G. Andersson, and K. Hook. Sentoy: A tangible interface to control the emotions of a synthetic character. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 1088–1089, 2003.
 13. G. Pfurtscheller and C. Neuper. Motor imagery and direct brain computer communication. *Proc. of the IEEE*, 89(7):1123–1134, July 2001.
 14. G. Pfurtscheller, R. Leeb, C. Keinrath, D. Friedman, C. Neuper, C. Guger, and M. Slater. Walking from thought. *Brain Research*, 1071:145–152, 2006.
 15. R. W. Picard. *Affective Computing*. MIT Press, 1997.
 16. H. Prendinger, C. Becker, and M. Ishizuka. A study in users' physiological response to an empathic interface agent. *Int'l J. Humanoid Robotics*, 3(3):371–391, 2006.
 17. H. Prendinger and M. Ishizuka. Human physiology as a basis for designing and evaluating affective communication with life-like characters. *IEICE Trans. Inf. & Syst.*, E88-D(11):2453–2460, 2005.
 18. H. Prendinger, J. Morib, and M. Ishizuka. Using human physiology to evaluate subtle expressivity of a virtual quizmaster in a mathematical game. *Int. J. Human-Computer Studies*, (62):231–245, 2005.
 19. M. S. Schwartz. *Biofeedback: A Practitioner's Guide*. New York: Guilford Press, 1995.
 20. A. Steed, J. Mortensen, and E. Frecon. Spelunking: Experiences using the DIVE system on CAVE-like platforms. In *Immersive Projection Technologies and Virtual Environments*, volume 2, pages 153–164. Springer-Verlag/Wien, 2001.
 21. P. Sundstrom, A. Stahl, and K. Hook. In situ informants exploring an emotional mobile messaging system in their everyday practice. *Int. J. of Human-Computer Studies*.
 22. P. Sundstrom, A. Stahl, and K. Hook. emoto: affectively involving both body and mind. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 2005–2008, 2005.
 23. N. Wiener. *Cybernetics, or Control and Communication in the Animal and the Machine*. MIT Press, 1961.