

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Modeling Forms of Surprise in an Artificial Agent

Permalink

<https://escholarship.org/uc/item/54p692s6>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Authors

Macedo, Luis
Cardoso, Amilcar

Publication Date

2001

Peer reviewed

Modeling Forms of Surprise in an Artificial Agent

Luís Macedo (lmacedo@isec.pt)

Instituto Superior de Engenharia de Coimbra / CISUC – Centro de Informática e Sistemas da Universidade de Coimbra, Quinta da Nora
3031-601 Coimbra, Portugal

Amílcar Cardoso (amilcar@dei.uc.pt)

Departamento de Engenharia Informática da Universidade de Coimbra / CISUC – Centro de Informática e Sistemas da Universidade de Coimbra, Pinhal de Marrocos
3030 Coimbra, Portugal

Abstract

Mainly rooted in the cognitive-psychoevolutionary model of surprise proposed by the research group of the University of Bielefeld (Meyer, Reisenzein, Schützwohl, etc.), the computational model of surprise described in this paper relies on the assumption that surprise-eliciting events initiate a series of mental processes that begin with the appraisal of unexpectedness, continue with the interruption of ongoing activity and the focusing of attention on the unexpected event, and end with the analysis and evaluation of that event plus revision of beliefs. With respect to the computation of unexpectedness, the model also incorporates suggestions by Ortony and Partridge. This model of surprise is implemented in an artificial agent called S-EUNE, whose task is to explore uncertain and unknown environments. The accuracy of our surprise model was evaluated in a series of experimental tests that focused on the comparison of surprise intensity values generated by the artificial agent with ratings by humans under similar circumstances.

Introduction

Roughly speaking, artificial and biological agents accept percepts from the environment and generate actions. Since different actions may lead to different states of the world, in order to perform well (to execute the “right” action), some kinds of artificial agents make use of a mathematical function that maps a state of the world onto a real number - the utility value. Thus, in those agents, decision-making is performed by selecting the action that leads to the state of the world with the highest utility (Russell & Norvig, 1995; Shafer & Pearl, 1990).

Although research in Artificial Intelligence has all but ignored the significant role of emotions in reasoning/decision-making (e.g., Damásio, 1994), several computational models for emotions have been proposed in the past years, based in part on research in psychology and neuroscience (for a detailed review of those models see e.g., Pfeifer, 1988; Picard, 1997).

Considered by many authors as a biologically fundamental emotion (e.g., Ekman, 1992; Izard, 1977), *surprise* may play an important role in cognitive activities, especially in attention focusing and learning (e.g., Izard, 1977; Meyer, Reisenzein, & Schützwohl, 1997; Ortony & Partridge, 1987; Reisenzein, 2000b) (note however, that some authors, like Ortony, Clore, and Collins, 1988, do not consider surprise an emotion). According to the research group of the University of Bielefeld, Germany (e.g., Meyer et al., 1997), surprise has two main functions, informational and motivational, that together promote both immediate adaptive actions to the surprising event and the prediction, control and effective dealings with future occurrences of the event. Ortony and Partridge’s view of surprise shares aspects with this model, especially in that both assume that surprise is elicited by unexpected events. The same is also true for Peters’ (1998) computational model of surprise, implemented in a computer vision system, that focuses on the detection of unexpected movements.

In this paper, we propose a computational model for surprise that is an adaptation (although with several simplifications) of the models proposed by the German research group of the University of Bielefeld and by Ortony and Partridge.

The following section presents an overview of the overall agent’s architecture into which the surprise model is integrated. Subsequently, we explain this model in detail. Finally, we describe experimental tests carried out to evaluate the accuracy of the surprise model.

Overview of the Agent’s Architecture

EUNE (Emotional-based Exploration of UNcertain and UNknown Environments) is an artificial agent whose goal is the exploration of uncertain and unknown environments comprising a variety of objects, and whose behavior is controlled by emotions, drives and other motivations. Besides desiring to know or be aware of the objects belonging to the environment, EUNE is also able to “feel” the emotions (including surprise)

those objects cause. In fact, these “felt emotions” guide the exploratory behavior of EUNE: roughly speaking, at any given time, among several objects available in the environment, EUNE selects that object for study and analysis that causes more positive emotion and less negative emotion (Izard, 1977) (see Reisenzein, 1996, for related theories of emotional action generation, and Barnes & Thagard, 1996, for an alternative approach to emotional decision-making). This process is repeated until all objects in the environment have become known.

In this article, we describe S-EUNE, a simplified version of EUNE whose emotional makeup is confined to the emotion of surprise. As many other agents, S-EUNE has perceptions, actions, goals, memory, emotions/drives, and deliberative reasoning/decision-making (Figure 1) (for more details on this architecture see Macedo & Cardoso, 2001).

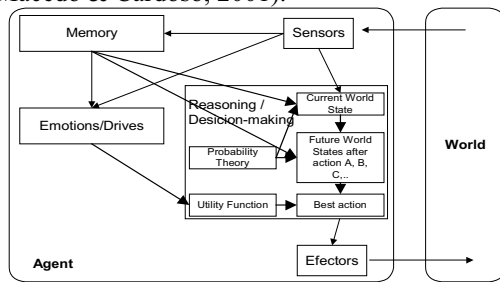


Figure 1: S-EUNE’s architecture.

Previously defined by the user, the environment comprises a variety of objects located at specific positions. In the present article, these objects are confined to buildings. Each object comprises three distinct, fundamental components: *structure*, *function* and *behavior* (Goel, 1992). For the sake of simplicity, the *structure* (the visible part of the object), is restricted to the shape of the object (e.g., triangular, rectangular, etc.); however, any object may comprise several sub-objects. The *function* of the object concerns its role in the environment (e.g., house, church, hotel, etc.). The *behavior* of the object concerns its activity (actions and reactions) in response to particular features of external or internal stimuli (e.g., static, mobile).

The perceptual system of the agent (two simulated sensors) provides information related to the *structure*, the *function*, and the *behavior* of the objects, as well as the distance of the objects. Note that the *function* of the objects is not accessible (i.e., cannot be inferred from visual information) unless the agent is at the same place as the object.

As a knowledge-based agent, S-EUNE stores all the information acquired through the sensors in its memory unit. The agent’s knowledge base is of an episodic kind: each object is stored, in the form of a graph, as a separate case in episodic memory. In addition, each

object representation is associated with a number that expresses its absolute frequency (Figure 2).





Field \ Case	C ₁	C ₂	C ₃	C ₄
Structure				
Function	House	House	Church	Hotel
Behavior	Static	Static	Static	Static
Abs. Freq.	50	40	5	5

Figure 2: Example of the episodic memory of S-EUNE after exploring an environment.

When information from the environment is sampled, the *surprise generation module* compares that information to the information stored in memory and outputs the intensity of the elicited surprise. A corresponding facial expression is also produced. The model of surprise is described in more detail in a later section. Note that, in the more general EUNE agent, this component - in this case called the *emotion, drives and other motivations module* - may also comprise other emotions apart from surprise, such as anger, sadness, etc., as well as drives and other motivations.

The *reasoning/decision-making module* of S-EUNE receives the information from the simulated external world and outputs the action that has been selected for execution. This module comprises several subprocesses: (i) taking the information of the world provided by the sensors (which may be incomplete) as input, the current state of the world (the agent’s current position, the position of the objects, etc.) is computed; (ii) taking the current state of the world, probability theory and memory-stored information as input, possible future world states and respective probabilities are computed for the actions that the agent can perform; (iii) from these actions, a single one (presumably the best one) is selected. This is the action that maximizes the Utility Function (Russell & Norvig, 1995; Shafer & Pearl, 1990), which in the case of S-EUNE relies heavily on the anticipated intensity of surprise elicited by the future state of the world. Thus, the preferences of S-EUNE are reflections of its anticipated surprise. In order to achieve this goal of maximizing anticipated surprise, the *reasoning/decision-making module* makes use of an Utility Function, abbreviated $U(W)$, which is based on the surprise function (defined in the next section) as follows:

$$U(W) = f(U_{surprise}(W)) = f(SURPRISE(Agt, Obj(W)))$$

This Utility Function means that the utility of a world state W is given by the surprise that the state W causes the agent to “feel”. In this article, a world state is defined as “being close to or seeing an object” (the object that is currently the focus of attention of the agent’s sensors), and f is taken to be the identity function, implying that $U(W)$ increases monotonically

with the intensity of surprise. As a consequence of this, the agent always selects for approach the object that actually elicits and/or promises to elicit maximum surprise.

Surprise Model

As mentioned before, our model of surprise is mainly based on Ortony and Partridge’s proposals and on the University of Bielefeld model. We will now give an overview of these models and then explain our computational model by comparing it with these two models.

Background Models

Ortony and Partridge (1987) proposed that there is a difference between surprisingness and expectation failure. They suggest that, although surprise sometimes results from expectation failure, it is frequently also caused by events for which expectations were never computed. In other words, one can be surprised by something one didn’t expect without having to expect something else. Ortony and Partridge also proposed that surprisingness is an important variable in artificial intelligence systems, particularly for attention and learning.

The following assumptions were made in their model: the system (or agent) receives an input proposition; the system has an episodic and semantic memory; elements of the memory may be immutable (propositions that are believed to be always true) or typical (those that are believed to be sometimes true); and, some elements of the memory are activated when an input proposition is received.

Ortony and Partridge further distinguish between *practically deducible propositions* and *practically non-deducible propositions*. *Practically deducible propositions* comprises the propositions that are explicitly represented in memory, as well as those that can be inferred from them by few and simple deductions. Hence, *practically deducible propositions* are that subset of formally deducible propositions that don’t require many and complex inferences. Furthermore, *practically deducible propositions* may be actively or passively deduced in a particular context. In the former case, their content corresponds to *actively expected* or *predicted* events; in the latter case, to *passively expected* (*assumed*) events.

Based on these assumptions, Ortony and Partridge proposed that surprise may result from three situations (Table 1 presents the correspondent range of values): (i) *active expectation failure*: here, surprise results from a conflict or inconsistency between the input proposition and an *active prediction* or *expectation*; (ii) *passive expectation failure* (or *assumption failure*): here, surprise results from a conflict or inconsistency between the input proposition and what the agent

implicitly knows or believes (*passive expectations* or *assumptions*); and (iii) *unanticipated incongruities* or deviations from norms: here, surprise results from a conflict or inconsistency between the input proposition (which in this case is a *practically non-deducible proposition*) and what, after the fact, may be judged to be normal or usual (cf. Kahneman & Miller, 1986), that is, *practically deducible propositions* (immutable or typical) that are suggested by the unexpected fact. Note that, in this case, at least prior to the unexpected event, there are no expectations (passive or active) with which the input proposition could conflict.

Table 1: Three different sources of surprise and correspondent range of values (adapted from Ortony & Partridge, 1987).

Confronted proposition	Related Cognition	
	Active	Passive
Immutable	[1]; $S_A=1$; <i>Prediction</i>	[2]; $S_P=1$; <i>Assumption</i>
Typical	[3]; $0 < S_A < 1$; <i>Prediction</i>	[4]; $S_P < S_A$; <i>Assumption</i>
Immutable	[5]; \emptyset	[6]; $S_P=1$; <i>none</i>
Typical	[7]; \emptyset	[8]; $0 < S_P < 1$; <i>none</i>

In their cognitive-psychoevolutionary model, the research group of the University of Bielefeld has made similar assumptions as Ortony and Partridge, namely that surprise (considered by them as an emotion) is elicited by the appraisal of unexpectedness. More precisely, it is proposed that surprise-eliciting events give rise to the following series of mental processes: (i) the appraisal of a cognized event as exceeding some threshold value of unexpectedness (schema-discrepancy) - according to Reisenzein (1999), this is achieved by a specialized comparator mechanism, the unexpectedness function, that computes the degree of discrepancy between “new” and “old” beliefs or schemas; (ii) interruption of ongoing information processing and reallocation of processing resources to the investigation of the unexpected event; (iii) analysis/evaluation of that event; (iv) possibly, immediate reactions to that event and/or updating or revision of the “old” schemas or beliefs.

Our Computational Model of Surprise

We have implemented a computational model of surprise, in the context of S-EUNE, that is an adaptation (although with some simplifications) of the University of Bielefeld’s model and in which the above-mentioned four mental processes elicited by surprising events are present. The suggestions by Ortony and Partridge are mainly concerned with the first of these steps, and are compatible with the Bielefeld model (see Reisenzein, 1999). Accordingly, we drew on these assumptions for the implementation of the appraisal of unexpectedness and the computation of the intensity of surprise, as well as the selection of knowledge structures in our model.

Within our model, knowledge is of an episodic kind, rather than being both semantic and episodic (although this will be part of our future work) as in Ortony and Partridge's model. Therefore, the knowledge structure of our model differs also from the schema-theoretic framework of the University of Bielefeld's model, that also assumes both episodic and semantic knowledge. In our model an input proposition (or new belief) is related to a visual object or parts of an object (for instance the visual effect of an object with squared windows, rectangular door, etc.). Besides, the agent has in its episodic memory explicit representations of similar propositions. Following Ortony and Partridge, we also distinguish between *deducible* and *non-deducible*, *active* and *passive*, *immutable* and *typical* propositions as well as between different possible sources of surprise (see Table 1). The immutability of a proposition can be extracted from the absolute frequency values associated with the cases (see Figure 2 above). For instance, the proposition "houses have squared facades" is immutable (since all the houses in memory have squared facades), whereas "houses have squared windows" is a typical proposition with a probability (immutability) value of .55 (as implied by Ortony and Partridge's model, in our model immutability is a continuous variable).

The usual activity of the agent is moving through the environment hoping to find buildings that deserve to be investigated. When one or more buildings are perceived, the agent computes expectations for their *functions* (for instance, "it is a house with 67% of probability", "it is a hotel with 45% of probability", etc.). Note that the *function* of a building is available to the agent only when its position and that of the building are the same. On the basis of this information (the *structure* of the object and predictions for its *function*), the agent then computes the surprise intensity that the building causes through the computation of its degree of unexpectedness (described below). Then, the building with the maximum estimated surprise is selected to be visited and investigated. This corresponds to the "interruption of ongoing activity" assumed in the Bielefeld model of surprise. The previously estimated value of surprise may now be updated with the additional information concerning the *function* of the building. The object is then stored in memory and the absolute frequencies of the affected episodes in memory are updated. This is a simplification of the fourth step of the University of Bielefeld's model (for alternative approaches to belief revision, see, for instance, Gärdenfors, 1988). Note that the experience of surprise is also accompanied by a correspondent facial expression (raised eyebrows, widened eyes, open mouth) (Ekman, 1992).

To see how the first step, the appraisal of unexpectedness, is performed, we now describe how the

degree of unexpectedness is computed in the three surprise-eliciting situations distinguished by Ortony and Partridge.

As said above, when the agent sees the structure of a building it computes expectations (*deducible*, *active expectations*) for its *function* (e.g., "it is a hotel with 45% of probability", etc.). If, after visiting that building, the agent finds out that it is a post office, it would be surprised, because its *active expectations* conflict with the input proposition (note that, in our model, belief conflicts may be partial rather as well as total). This is thus an example of the first source of surprise distinguished by Ortony and Partridge. In contrast, when the agent sees a building with a window (or roof, etc.) of a particular shape (for instance, circular), although it may not have made an *active prediction* for its shape, it is able to infer that it expected a rectangular shape with, for instance, 45% probability, a squared shape with 67%, etc. This is an example of a *deducible*, *passive expectation*: although not made before the agent perceived the building, it could easily infer an expectation for the shape of the window after it was perceived. This case is therefore an example of the second source of surprise because the input proposition "has a circular window" conflicts with the agent's *passive expectations*. Finally, when the agent sees a building with no facade, it has neither an *active* nor a *passive expectation* available, because there are no buildings with no facade in its memory and therefore the agent could not predict that. Thus, "the house has no facade" is an example of a *non-deducible proposition*. This is an example of the third source of surprise: there is a conflict between the input proposition "the house has no facade" and what after the fact is judged to be normal or usual ("buildings have a facade").

Let us now describe how the intensity of surprise is computed. There is experimental evidence supporting that the intensity of felt surprise increases monotonically, and is closely correlated with the degree of unexpectedness (see Reisenzein, 2000b, for a review of these experiments). This suggests that unexpectedness is the proximate cognitive cause of the surprise experience. On the basis of this evidence, we propose that the surprise felt by an agent *Agt* elicited by an object *Obj_k* is proportional to the degree of unexpectedness of *Obj_k*, considering the set of objects present in the memory of the agent. According to probability theory (e.g., Shafer & Pearl, 1990), the degree of expecting that an event X occurs is given by its probability P(X). Accordingly, the improbability of X, denoted by 1-P(X), defines the degree of not expecting X, and the intensity of surprise can, for simplicity, be equated with unexpectedness:

$$SURPRISE(Agt, Obj_k) = DegreeOfUnexpectedness(Obj_k, Agt(Memory)) = 1 - P(Obj_k)$$

Although other probabilistic methods might be used to compute $P(X)$, in the case of objects comprising several components we propose to compute the probability of the whole object Obj_k as the mean of the conditional probabilities of their n constituent parts, which are individually computed using Bayes's formula (Shafer & Pearl, 1990) (note that each one of those conditional probabilities individually gives the degree of unexpectedness of a specific piece of the object, given as evidence the rest of the object):

$$P(Obj_k) = \frac{\sum_{l=1}^n P(Obj_k^l | Obj_k^1, Obj_k^2, \dots, Obj_k^{l-1}, Obj_k^{l+1}, \dots, Obj_k^n)}{n}$$

Experimental Tests

Although our model is consistent with the experimental evidence reported, we performed two new experiments to test the following issues: (i) whether the intensity values generated by the artificial agent match those of humans under similar circumstances; (ii) the role of the amount of previous knowledge on the surprise intensity; (iii) whether the surprise intensity values generated by the artificial agent fall within the range of the surprise intensity values proposed in Ortony and Partridge's model. In both experiments, the participants (S-EUNE and 60 humans with mean age of 20.5 years) were presented with 40 quiz-like items. Experiment 1 was performed in an abstract domain with hedonically neutral events (see Stiensmeier-Pelster, Martini, & Reisenzein, 1995, for a similar experiment with humans). Each "quiz item" consisted of several sequences of symbols. Some of the "quiz items" contained a single sequence in which one symbol was missing. Experiment 2 was performed in the domain of buildings. In this case, each "quiz item" consisted of the presentation of a building, and some items did not include information about its *function* (see Reisenzein, 2000a, for a conceptually similar experiment with humans). In those cases where a symbol of the sequence (Experiment 1) or information about the *function* of the building (Experiment 2) was missing, the participants had to state their expectations for the missing symbol or the missing *function*. Subsequently, the "solution" (the missing information) of the "quiz item" was presented and the participants were asked to rate the intensity of felt surprise about the "solution", as well as for the whole sequence/building. For "quiz items" ending with complete sequences or complete buildings, the participants had to rate the intensity of felt surprise about a specified element of the sequence or a specified piece of the building. Subsequently, they also indicated their *passive expectations* for that element/piece. The "quiz items" used in both experiments were selected on the basis of a previous questionnaire study. They were equally distributed

among the three sources of surprise described earlier, as well as among different intensities of surprise ranging from low to high.

Figure 3 presents the results of Experiment 1. It can be seen that the intensity of surprise computed for an element of a sequence by the agent (labeled *S-EUNE-Piece* in Figure 3) is close (average difference = .065, i.e., 6.5%) to the corresponding average intensity given by the human judges (*Humans Average-Piece*). Even better results (average difference = .022) were obtained for the surprise values computed for the whole sequence (*S-EUNE-Whole* and *Humans Average-Whole*). Figure 3 also shows that the standard deviations of the surprise intensities given by the 60 humans (*S.D.-Humans-Piece*, *S.D.-Humans-Whole*) were less than .23 (for an element) and .18 (for the whole sequence).

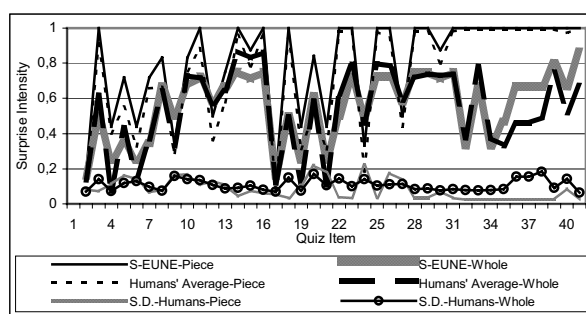


Figure 3: Results of Experiment 1.

Figure 4 presents the results of Experiment 2. In this experiment, S-EUNE answered the "quiz items" several times, each time with a different episodic memory. Due to the lack of space, we reported only the results of three sessions, denoted by S-EUNE-I, IV and V (with I, IV and V denoting an increasingly large memory). It can be seen that the surprise values of the agent are not as close to the human judgments as in the previous domain. For instance, the average differences for S-EUNE-V were .47 (for a piece of a building) and .05 (for the whole building). This happened most likely because, in contrast to the previous, hedonically neutral domain, in the domain of buildings the knowledge of humans and of the agent is different. However, the results suggest that the larger the episodic memory, and the closer its probability distribution corresponds to the real world, the closer are the surprise values given by the agent and by the humans. For instance, S-EUNE-V (*S-EUNE-V-Piece* and *S-EUNE-V-Whole*) showed the best correspondence to the human ratings. This experiment also confirms to some extent the dependence of surprise on the contents and developmental stage of memory, suggested by studies that compared the surprise reactions of adults with those of children (Schützwohl & Reisenzein, 1999).

Both experiments also confirmed that the values of surprise fall in the ranges predicted by Ortony and Partridge, with the exception that, in the case of the source of surprise corresponding to cell [8] of Table 1, the values are always 1, and, in the case of cell [4], $S_p=S_A$.

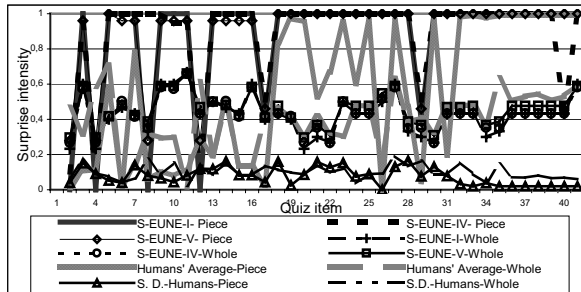


Figure 4: Results of Experiment 2.

Conclusions

The results of the reported experiments suggest that the described computational model is a possible model of surprise. However, alternative surprise functions are conceivable, such as, $SURP(O)=\ln_2(1/P(O))$ (as suggested by information theoretic accounts) or $SURP(O)=1-P(O)\Leftarrow P(O)<.5$; $SURP(O)=0\Leftarrow P(O)\ge.5$ (as suggested to us by Rainer Reisenzein). We are currently exploring these and other alternatives.

Acknowledgments

We would like to thank Andrew Ortony and Rainer Reisenzein for their helpful comments, and the participants of our experiments for their cooperation.

References

Barnes, A., & Thagard, P. (1996). Emotional decisions. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 426-429). Mahwah, NJ: Erlbaum.

Damásio, A. (1994). *Descartes' error, Emotion Reason and the Human Brain*. New York: Grosset/Putnam Books.

Ekman, P. (1992). An argument for basic emotions. In N. L. Stein & K. Oatley (Eds.), *Basic Emotions* (pp. 169-200). Hove, UK: Erlbaum.

Gärdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states*. Cambridge, MA: Bradford Books.

Goel, A. (1992). Representation of design functions in experience-based design. In D. Brown, M. Waldron, & H. Yoshikawa (Eds.), *Intelligent Computer Aided Design* (pp. 283-308). Amsterdam: North-Holland.

Izard, C. (1977). *Human Emotions*. New York: Plenum.

Kahneman, D., & Miller, D. T. (1986). Norm theory: comparing reality to its alternatives. *Psychological Review*, 93, 136-153.

Macedo, L., & Cardoso, A. (2001). SC-EUNE – Surprise/Curiosity Exploration of UNcertain and UNKnown Environments. *Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing* (pp. 73-81). York, UK: SSAISB.

Meyer, W., Reisenzein, R., & Schützwohl, A. (1997). Towards a process analysis of emotions: The case of surprise. *Motivation and Emotion*, 21, 251-274.

Ortony, A., & Partridge, D. (1987). Surprisingness and Expectation Failure: What's the Difference?. *Proceedings of the 10th International Joint Conference on Artificial Intelligence* (pp. 106-108). Los Altos, CA: Morgan Kaufmann.

Ortony, A., Clore, G., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.

Peters, M. (1998). Towards artificial forms of intelligence, creativity, and surprise. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 836-841). Mahwah, NJ: Erlbaum.

Pfeifer, R. (1988). Artificial intelligence models of emotion. In V. Hamilton, G. Bower & N. Frijda (Eds.), *Cognitive Perspectives of Emotion and Motivation* (pp. 287-320). London: Kluwer.

Picard, R. (1997). *Affective Computing*. Cambridge, MA: MIT Press.

Reisenzein, R. (2000a). Exploring the strength of association between the components of emotion syndromes: The case of surprise. *Cognition and Emotion*, 14, 1-38.

Reisenzein, R. (2000b). The subjective experience of surprise. In H. Bless & J. Forgas (Eds.), *The message within: The role of subjective experience in social cognition and behavior* (pp. 262-279). Philadelphia, PA: Psychology Press.

Reisenzein, R. (1999). A theory of emotions as metarepresentational states of mind. *Personality and Social Psychology Review*. (Under review)

Reisenzein, R. (1996). Emotional action generation. In W. Battmann & S. Dutke (Eds.), *Processes of the molar regulation of behavior* (pp. 151-165). Lengerich: Pabst Science Publishers.

Russell, S., & Norvig, P. (1995). *Artificial Intelligence - A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.

Shafer, G., & Pearl, J. (Eds.) (1990). *Readings in Uncertain Reasoning*. Palo Alto, CA: Morgan Kaufmann.

Schützwohl, A., & Reisenzein, R. (1999). Children's and adults' reactions to a schema-discrepant event: A developmental analysis of surprise. *International Journal of Behavioral Development*, 23, 37-62.

Stiensmeier-Pelster, J., Martini, A., & Reisenzein, R. (1995). The role of surprise in the attribution process. *Cognition and Emotion*, 9, 5-31.