A Surprise-based Agent Architecture

Luis Macedo, Amilcar Cardoso
Centre for Informatics and Systems of the University of Coimbra
Department of Informatics, Polo II
3030 Coimbra, Portugal
{macedo,amilcar}@dei.uc.pt

Rainer Reisenzein
Institute for Psychology, University of Greifswald
Department of General Psychology II
Franz-Mehringstr. 47
17487 Greifswald, Germany
rainer.reisenzein@uni-greifswald.de

Abstract

We describe the architecture of an artificial agent whose task is to explore unknown environments. It is assumed that the behavior of the agent is partly (and in a simplified version of the agent, even wholly) controlled by surprise. We describe the computational model of surprise incorporated in the architecture, briefly report on related empirical research on surprise intensity in humans, and summarize the results of simulation studies that compared a purely surprise-motivated agent to agents with additional motives.

1 Introduction

Considered by many authors to be a biologically fundamental emotion (e.g., Ekman, 1992; Izard, 1991), surprise may play an important role in the cognitive activities of intelligent agents, especially in attention focusing (Izard, 1991; Meyer, Reisenzein, & Schützwohl, 1997; Ortony & Partridge, 1987; Reisenzein, 2000), learning (Schank, 1986) and creativity (Boden, 1995; Williams, 1996). Psychological experiments (e.g., Meyer et al., 1997) provide evidence that surprising-eliciting events initiate a series of mental processes that (a) begin with the appraisal of a cognized event as exceeding some threshold value of unexpectedness or schema discrepancy, (b) continue with the interruption of ongoing information processing and the reallocation of processing resources to the surprise-eliciting event, and (c) culminate in the analysis and evaluation of that event plus immediate reactions to it and/or schema (belief) updating/revision. According to Meyer et al., surprise has two main functions, the one informational and the other motivational: it informs the individual about the occurrence of a schema-discrepancy, and it provides an initial impetus for the exploration of the unexpected event. Thereby, surprise promotes both immediate adaptive actions to the unexpected event and the prediction, control and effective dealings with future occurrences of the event.

Ortony and Partridge's (1987) model of surprise shares several aspects with the one proposed by Meyer et al. (1997), especially the assumption 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. Finally, models of surprise have also been proposed in the fields of knowledge discovery and data mining (e.g., (Suzuki & Kodratoff, 1998)).

Inspired by the models of Meyer et al. (1997) and Ortony and Partridge (1987), Macedo and Cardoso (2001) developed a computational model of surprise in the context of a more general agent architecture. This computational model was elaborated further by Macedo, Reisenzein, and Cardoso (2004), who discussed different possible functions for the computation of surprise intensity and evaluated these functions in an empirical study. Additional empirical research of the same nature was conducted by Reisenzein and Macedo (2006). In addition, the agent architecture has been further refined (e.g., by adding a semantic memory component to the existing episodic memory).

The following section presents an overview of the overall architecture of the agent, into which the surprise model is integrated. Subsequently, we explain the computational model of surprise in more detail. Finally, we briefly report about results of our empirical research on surprise intensity in humans and summarize the results of simulation studies that compared a purely surprise-motivated agent to agents with additional motives.

2 Agent Architecture

EUNE (Emotion-based Exploration of UNcertain and UNknown Environments) is an artificial agent whose goal is the exploration of unknown environments comprising a variety of objects, and whose behavior is controlled by emotions, drives and other motivations (for more detail, see Macedo & Cardoso, 2004). Besides desiring to get to know the objects in its environment, EUNE also "feels" the emotions (including surprise) those objects cause. In fact, currently "felt" as well as anticipated 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 maximizes desired emotions (Izard, 1977) (see Reisenzein, 1996), for related and alternative theories of emotional action generation). This process is repeated until all objects in the environment have become known to the agent.

In this article, we describe S-EUNE, a simplified version of EUNE whose emotional makeup is confined to the emotion of surprise. The architecture of S-EUNE, like that of EUNE (Figure 1) is based on the BDI approach (Bratman, Israel, & Pollack, 1988). Similar to many other agent architectures, it includes the following modules: sensors/perception; effectors/actuators; memory/beliefs; emotions, drives and other motivations; intentions/goals; desires; and, deliberative reasoning/decision-making. The deliberative reasoning/decision-making module is at the core of the architecture. It receives internal information (from memory) and environmental information (through the sensors) and outputs an action that has been selected for execution. The process of action selection takes into account the states of the environment the agent would like to happen (desires). The agent's preferences are represented implicitly by means of a mathematical function that evaluates states of the world in terms of their utility for the agent. Based on this function, the decision-making module selects the action that maximizes utility (Russell & Norvig, 1995). In the case of EUNE, the agent's utility is assumed to reside in (or to be derived from) actual as well as anticipated emotional feelings (which, in the case of S-EUNE, are restricted to surprise). The intensities of these feelings are computed by the emotion/motivation module, taking into account both the past experience of the agent (the information stored in memory) and information about the current environment provided by the sensors.

The simulation environment used as a test bed for our approach to exploration comprises a variety of entities located at specific positions. In the example used below to illustrate the operation of S-EUNE, these entities are confined to buildings characterized by structural and functional properties; but in principle all kinds of entities, including other animated agents, can be simulated. The *structure* of a building comprises the shape (triangular, rectangular, etc.) of the roof, facade, door and windows. The possible *functions* of a building include: house, church, hotel, and hospital.

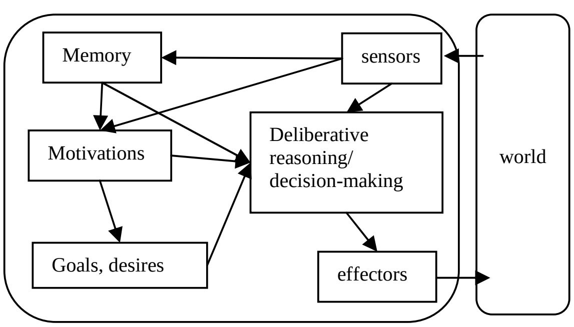

Figure 1 - Agent’s architecture.

Being a knowledge-based agent, S-EUNE stores all information acquired through the sensors in its memory unit. This includes information about the composition of the agent's environment, such as the position of the entities (objects and other animated agents) that inhabit it, the structural and functional properties of these entities, and the actions executed by other agents. This information is stored in several different memory units. A grid-based metric map (Thrun, 2002) is used to model the spatial structure of the agent's physical environment. Descriptions of entities (concerning their physical structure and function and in the case of other agents, their actions) are stored in both an episodic and a semantic memory unit (Tulving, 1972). The physical structure of an entity can be described analogically or propositionally (Eysenck & Keane, 1991). The function of an entity is described by specifying the entity's role in the environment (e.g., "is a house", "is a car", "is an agent"). Both the description of the physical structure and the description of the function of an entity can be probabilistic (e.g., "is a house with probability .70"). Concrete entities (i.e., entities represented in episodic memory) with similar features may be generalized or abstracted into an abstract entity or prototype, which is stored in the *semantic memory for entities* (for further description, see Macedo & Cardoso, 2004).

Whenever information from the environment is sampled, the *surprise generation module* compares this information to the information stored in memory and outputs a signal that reflects the intensity of surprise elicited by the new information. This module is described in more detail in later section. Here, it is important to note that, although S-EUNE’s emotions are deliberately confined to surprise, the more general agent EUNE may also have other emotions (e.g., fear) and motivations (e.g., curiosity or hunger) in addition to surprise.

To explore its environment, the agent is continuously engaged in deliberative reasoning/decision-making. More precisely speaking, at any given time, the agent senses the environment and computes the current world state (the location, structure and function of the entities in its surroundings), based on both sensorial information and on the generation of expectations for missing pieces of information. On the basis of the computed world state *W*, a goal of the type *visitEntity* is generated for each entity that has not yet been visited. In addition, a goal of type *visitLoc* is generated for all frontier cells, that is, the cells on the boundary of the currently known part of the agent's environment (Simmons et al., 2000). These goals are inserted into a ranked list of goals which may already contain goals that were generated in the past but have not yet been accomplished. The top goal of this list is the one that that maximizes the utility function of the agent (Russell & Norvig, 1995; Shafer & Pearl, 1990). In the special case of S-EUNE, this utility is exclusively based on the intensity of surprise elicited by the state of the world *W*. That is, in contrast to humans and to the more general agent EUNE, S-EUNE's only desire is to explore objects that elicit surprise.

To this end, the *reasoning/decision-making module* uses an utility function *EU(W)* that is based on the surprise function (defined in the next section) as follows:

$$EU(W) = f(U_{surprise}(W)) = f(SURPRISE(W))$$

That is, the utility of world state W is a function of the surprise elicited by W. In this article, a world state is defined as "seeing an object" (the object that is currently at the focus of attention of the agent's sensors), and f is taken to be the identity function, implying that $EU(W)$ increases monotonically with the intensity of surprise. As a consequence, the agent always selects for approach the object that elicits maximum surprise.

3 The Surprise Module of S-EUNE

The surprise module of S-EUNE (Macedo & Cardoso, 2001; Macedo et al., 2004) is mainly based on Ortony and Partridge's (1987) proposals and on those of Meyer et al. (1997). Therefore, we first briefly review these background theories.

3.1 Background Models

In line with most other surprise theorists, Ortony and Partridge (1987) assume that surprise is caused by events that are commonsensically called unexpected. However, different from other theorists, they proposed that unexpectedness, and hence surprise, covers two different cases. First, surprise is elicited when prior expectations regarding an event are disconfirmed. In addition, however, surprise can also be elicited by some events for which one had no explicit expectations (neither conscious nor unconscious).

In more detail, similar to S-EUNE, Ortony and Partridge's model of surprise assumes a system (or agent) with an episodic and a semantic propositional memory whose elements may be immutable (propositions that are believed to be always true) or typical (propositions that are believed to be usually but not always true). Furthermore, Ortony and Partridge distinguish between practically deducible and practically non-deducible propositions. *Practically deducible propositions* comprise all propositions that are explicitly represented in memory, as well as those that can be inferred from these by few and simple deductions. Finally, practically deducible propositions may be either actively or passively deduced. 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 results when the system encounters a conflict or inconsistency between an input proposition and (a) preexisting representations or (b) representations computed "after the fact". More precisely, surprise results in each of three situations (Table 1 shows the corresponding 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*). (iii) *Unanticipated incongruities* or deviations from norms: here, surprise results from a conflict or inconsistency between the input proposition and what, after the fact, is judged as normal or usual (Kahneman & Miller, 1986); that is, between the input proposition and practically deducible propositions (immutable or typical) that are suggested by the unexpected fact. Note that, in this case, prior to the unexpected event there are no explicit expectations (passive or active) with which the input proposition could conflict.

Table 1: Three different sources of surprise and corresponding value ranges (adapted from (Ortony & Partridge, 1987)).

Confronted proposition	Related Cognition	
	Active	Passive
Immutable	[1]; S_A=1; *Prediction*	[2]; S_P=1; *Assumption*
Typical	[3]; 0< S_A<1; *Prediction*	[4]; S_P<S_A; *Assumption*
Immutable	[5]; ∅	[6]; S_P=1; *none*
Typical	[7]; ∅	[8]; 0< S_P<1; *none*

The cognitive-psychoevolutionary model of surprise proposed by Meyer et al. (1997) also assumes that surprise is elicited by the appraisal of unexpectedness. However, this model also makes assumptions about the cognitive and behavioral *consequences* of this appraisal. In more detail, the authors assume that incoming information is continuously compared with preexisting activated schemas or expectations. According to Reisenzein (2001), this is achieved by a specialized, hardwired comparator mechanism that preconsciously computes the degree of discrepancy between "new" and "old" beliefs or schemas. If the degree of unexpectedness or schema-discrepancy exceeds some threshold value, surprise is felt, ongoing information-processing is disrupted, and processing resources are reallocated to the investigation of the unexpected event. This investigation typically comprises the analysis of the causes of the event and its significance for the person's goals. Finally, the person's schemas or beliefs may be updated, and adaptive reactions to the event may be executed.

3.2 Overview of the Computational Model of Surprise

In line with the surprise model of Meyer et al. (1997), in S-EUNE input propositions (or newly acquired beliefs) about objects or events (e.g., the belief that an object with square windows is located at a certain position) are continuously compared with existing representations of objects or events in memory. 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 value of a proposition can be extracted from the absolute frequency values associated with the cases stored in episodic memory. For instance, if all houses represented in episodic memory have a pentagonal shape, the proposition "houses have a pentagonal shape" is immutable; whereas if only half of the houses have square windows, the proposition "houses have square windows" is a typical proposition with a probability (immutability) value of 0.50.

The usual activity of the agent consists of moving through the environment "hoping" to find interesting things (objects or events) that deserve to be investigated. Although we assume that exploratory behavior in humans can be (and probably typically is) in the service of several different motives, in S-EUNE, we deliberately ignored all possible motives of exploration except the motive to investigate surprising objects, to be able to study the behavior of a purely surprise-motivated agent.

When one or more objects or events are perceived by the agent, it computes expectations for missing (i.e. not directly perceivable) pieces of information. In the present implementation, for example, the agent is able to perceive part of the structural properties of buildings within its range of vision with certainty, but it does not know the function of a building with certainty until its position and that of the building are identical. However, even prior to this point, the agent can still form expectations about the object's function (e.g., "it is a house with probability .70", "it is a hotel with probability .45"). On the basis of the available information (e.g., the visible *structure* of an object) and the computed expectations (e.g., predictions of the *function* of an object), the agent then computes the intensity of surprise caused by the object. These computations, which correspond to the "appraisal of unexpectedness" in the surprise model of Meyer et al. (1997), are described in more detail below. Subsequently, the object with the maximum surprise intensity value is selected to be visited and investigated. This corresponds to the "interruption of ongoing activity" and the "reallocation of processing resources" assumed in the Meyer er al. model. On the basis of the additional information acquired about the object or event, the surprise intensity value may be updated. Finally, the object/event is stored in memory and the absolute frequencies of the affected objects/events in memory are updated. This is a simplification of the fourth step of the Meyer, Reisenzein and Schützwohl model (for alternative approaches to belief revision, see e.g. (Gärdenfors, 1988)).

The different surprise-eliciting situations distinguished by Ortony and Partridge are dealt with in S-EUNE in the following way. As said above, when the agent perceives an object, it first computes expectations (*deducible*, *active expectations*) for missing information (e.g., "it is a hotel with probability .45"). If, after having visited that object, the agent detects that the object is different from what was expected (e.g., if it is a post office), the agent is surprised because its *active expectations* conflict with the input proposition. This is thus an example of the first source of surprise distinguished by Ortony and Partridge. In contrast, when S-EUNE perceives an object with particular properties (e.g., a building with a window of circular shape) that were not actively predicted, it may still be able to infer that it expected something (e.g., a rectangular window with probability .40, a square window with probability .67, etc.). This is an example of a *deducible*, *passive expectation*: although the expectation was not present before the agent perceived the object, it was inferred after the object had been perceived. This case is therefore an example of the second source of surprise distinguished by Ortony and Partridge, where an input proposition conflicts with an agent's *passive expectations*. Finally, when an agent perceives an object with a completely new part (e.g., a building with no facade), it has neither an active nor a passive expectation available. The reason is that, because there are no objects of this kind (e.g., buildings with no facade) stored in the agent's memory, the agent cannot predict that such objects might be encountered. The perception of an object with a completely new part is thus an example of a *non-deducible proposition*. This is an example of the third source of surprise distinguished by Ortony and Partridge: there is a conflict between the input proposition (e.g., "the house has no facade") and what *after the fact* is judged to be normal or usual (e.g., "buildings have a facade").

Let us now address the question of how the intensity of surprise should be computed. In humans, this problem has already been solved by evolution; therefore, a reasonable approach is to model the agent's surprise function according to that of humans. Experimental evidence from human participants summarized in (Reisenzein, 2000) suggests that the intensity of felt surprise about an event E increases monotonically, and is closely correlated with, the degree of unexpectedness of E, defined as $1\text{-}P(E)$, with $P(E)$ being the subjective probability of E. On the basis of this evidence, we assume that the surprise "felt" by an agent about an event E is an (at least weakly) monotonically increasing function of $1\text{-}P(E)$ (Macedo & Cardoso, 2001). To determine the shape of the surprise function more precisely, Macedo et al. (2004) compared several theoretically derived surprise functions with surprise intensity ratings of human participants in the domain of political elections and sports games. The results of this study suggested: (a) the intensity of experienced surprise may not only depend on the probability of the actual outcome, $P(E)$, but also on the probability of alternative outcomes, in particular on that of the maximally probable event of a set of mutually exclusive alternatives; and (b) the shape of the surprise function may be nonlinear. Specifically, Macedo et al. (2004) found that of the surprise functions considered, the one that best fit the data was the following:

$$SURPRISE(E) = \log_2(1 + P(E_{\max}) - P(E))$$

In this formula, E_{max} is the event with the *highest* probability in the set of possible outcomes. This formula implies (a) within each set of mutually exclusive events, there is always at least one (E_{max}) whose occurrence is unsurprising, namely the event with the maximum probability in the set; and (b) for the remaining events E in the set, the surprise caused by their occurrence is proportional to the logarithm of the difference between $P(E_{max})$ and their probability $P(E)$.

This difference can be interpreted as the amount by which P(E) has to be increased for E to become unsurprising. The surprise function also predicts that maximum surprise, namely SURPRISE(E) = 1, occurs only if P(E_{max}) = 1 and hence P(E) = 0. The logarithmic transformation captures the idea that surprise function is nonlinear.

Following the initial research by Macedo et al. (2004), Reisenzein and Macedo (2006) conducted a more extensive set of experiments that focused on surprise elicited by (more or less) unexpected quiz solutions and by unexpected gains and losses in monetary lotteries. The data analysis of these studies is still under way; however, preliminary results provide only partial support for the conclusions drawn by Macedo et al. (2004). Specifically, although there was again evidence that the intensity of surprise in humans depends not only of the probability of the actual outcome, but also on that of alternative outcomes (see also Teigen & Keren, 2003), the effect of alternative outcomes was comparatively weak and appeared to be restricted to particular situations (selected findings from this study will be presented at the conference).

Finally, note that the above equation only describes the intensity of surprise elicited by a single input proposition (e.g., "this object has round windows"), whereas S-EUNE computes the intensity of surprise elicited by complete objects (e.g., a building with round windows and a square door that is a church). To compute the total surprise elicited by an object, the surprise function is applied to the components of the object representation (e.g., the individual components of the complete propositional description of the object), taking into account their probabilistic dependencies (for more detail, see Macedo & Cardoso, 2004).

4 Conclusions

S-EUNE is based on the assumption that surprise has informational value that is useful for exploration. As explained, S-EUNE's exploration strategy leads the agent to focus on those aspects or parts of its environment that elicit the most surprise. Surprise signals to the agent that one or several of its assumptions about the world were false and need to be modified. However, as argued by Meyer et al. (1997), adequate belief updating frequently requires further exploration of the surprising event (e.g., concerning its causes).

In the current version of S-EUNE, surprise influences exploration by having a *direct*, monotonic effect on the utility function of the agent (the utilities of different exploration goals). However, it may be instructive to study alternative versions of S-EUNE in which the effect of surprise on goal selection is more indirect. Indeed, we believe that one of the main benefits of agent research for psychology is the possibility to comparatively study the behavior of agents with different but related motivational and emotional mechanism. For example, one could assume that exploratory behavior is directly motivated by curiosity, but that surprise is one of the factors that elicit curiosity (e.g., Berlyne, 1971). Relatedly, one could assume that surprise-based exploration is a form of sensation-seeking (e.g., Zuckerman, 1979). However, it should be acknowledged that this form of sensation-seeking is of a highly specific, epistemic kind.

S-EUNE was deliberately constructed as an agent motivated exclusively by the desire to explore surprising objects and events. This was done to be able to study surprise-motivated behavior in "pure" form. It was not meant to imply that realistic agents (including humans) are only motivated by this goal. In fact, simulation studies found that, when judged by the ability to explore an environment *efficiently*, S-EUNE performed worse than an agent that used a more traditional search strategy (e.g., minimizing the distance traversed, or maximizing the amount of information expected to be acquired). Additional simulation experiments revealed that the agent explored the entities in its environment more quickly if it used a strategy that took into account hunger, either alone or combined with surprise. Whereas an exclusively surprise-based exploration strategy lead to erratic search paths, much more orderly exploration paths were obtained if additional motives, such as hunger, were taken into account. A main reason for this was that the agent's tendency to visit distant, surprising entities or locations was considerably restrained if the agent also expected to experience hunger at these distant locations. On the other hand, a surprise-based exploration strategy proved to be useful when exploration was performed in the context of creativity, that is, when the primary goal of the agent was to encounter artistically or scientifically creative entities. This strategy could therefore be useful, for example, for agents who explore environments such as museums, or who search the scientific literature for unusual, interesting ideas. Indeed, (Davis, 1971) argued that one of the main factors that determine interest in a scientific proposition is its deviation from the beliefs or expectations of the audience.

References

Berlyne, D. E. (1971). *Aesthetics and Psychobiology*. New York: Appleton-Century-Crofts.

Boden, M. (1995). Creativity and unpredictability. *SEHR, 4*(2).

Bratman, M., Israel, D., & Pollack, M. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence, 4,* 349–355.

Davis, M. (1971). That's Interesting! Towards a Phenomenology of Sociology and a Sociology of Phenomenology. *Phil. Soc. Sci., 1,* 309-344.

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

Eysenck, M., & Keane, M. (1991). *Cognitive psychology*. London: Lawrence Erlbaum Associates.

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

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

Izard, C. (1991). *The Psychology of Emotions*. New York: Plenum Press.

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

Macedo, L., & Cardoso, A. (2001). Modelling Forms of Surprise in an Artificial Agent. In J. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 588-593). Mahwah, NJ: Erlbaum.

Macedo, L., & Cardoso, A. (2004). Exploration of Unknown Environments with Motivational Agents. In N. Jennings & M. Tambe (Eds.), *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 328 - 335). New York: IEEE Computer Society.

Macedo, L., Reisenzein, R., & Cardoso, A. (2004). Modeling Forms of Surprise in Artificial Agents: Empirical and Theoretical Study of Surprise Functions. In K. Forbus & D. Gentner & T. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

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.

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.

Reisenzein, R. (1996). Emotional Action Generation. In W. Battmann & S. Dutke (Eds.), *Processes of the molar regulation of behavior*. Lengerich: Pabst Science Publishers.

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

Reisenzein, R. (2001). Appraisal processes conceptualized from a schema-theoretic perspective: Contributions to a process analysis of emotions. In K. Scherer & A. Schorr & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, Methods, Research* (pp. 187-201). Oxford: Oxford University Press.

Reisenzein, R., & Macedo, L. (2006). Surprise and subjective probability. Unpublished.

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

Schank, R. (1986). *Explanation Patterns: Understanding Mechanicaly and Creatively*. Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun, S., & Younes, H. (2000). Coordination for Multi-Robot Exploration and Mapping, *Proceedings of the AAAI-2000*.

Suzuki, E., & Kodratoff, Y. (1998). Discovery of Surprising Exception Rules Based on Intensity of Implication. In J. Zytkow & M. Quafafou (Eds.), *Proceedings of Second European Symposium on Principles of Data Mining and Knowledge Discovery, PKDD '98* (pp. 10-18). Berlin: Springer.

Teigen, K. H., & Keren, G. (2003). Surprises: Low probabilities or high contrasts? *Cognition and Emotion, 87*, 55-71.

Thrun, S. (2002). Robotic mapping: A survey. In G. Lakemeyer & B. Nebel (Eds.), *Exploring Artificial Intelligence in the New Millenium*. San Mateo, CA: Morgan Kaufmann.

Tulving, E. (1972). Episodic and Semantic Memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381-403). New York: Academic Press.

Williams, M. (1996). Aesthetics and the explication of surprise. *Languages of Design, 3*, 145-157.

Zuckerman, M. (1979). *Sensation seeking: Beyond the optimal level of arousal*. Hillsdale, NJ: Lawrence Erlbaum Associates.