Modeling Forms of Surprise in Artificial Agents: Empirical and Theoretical Study of Surprise Functions

Luís Macedo (lmacedo@isec.pt)
Department of Informatics and Systems Engineering, Engineering Institute, Polytechnic Institute of Coimbra, Quinta da Nora
3030-199 Coimbra, Portugal
Centre for Informatics and Systems of the University of Coimbra, Department of Informatics, Polo II
3030 Coimbra, Portugal

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

Amílcar Cardoso (amilcar@dei.uc.pt)
Centre for Informatics and Systems of the University of Coimbra, Department of Informatics, Polo II
3030 Coimbra, Portugal

Abstract

This paper addresses the issue of how to compute the intensity of surprise in an artificial agent. Resolution of this issue is important for the further specification of the computational model of surprise proposed by Macedo and Cardoso (2001) that was implemented in artificial agents "living" in a multi-agent environment. This model of surprise is mainly rooted in the cognitive-psychoevolutionary model of surprise proposed by the research group of the University of Bielefeld (Meyer, Reisenzein, & Schützwohl, 1997) and in proposals by Ortony and Partridge. We propose several possible functions to compute the intensity of surprise. To assess their accuracy, they were evaluated in an experimental test that focused on the comparison of surprise intensity values generated by artificial agents with ratings by humans under similar circumstances.

Introduction

Considered by many authors 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 et al., 1997; Ortony & Partridge, 1987; Reisenzein, 2000b), learning (Schank, 1986) and creativity (Boden, 1995; Williams, 1996). Psychological experiments conducted by Meyer, Reisenzein and Schützwohl 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 these authors, 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, Reisenzein and Schützwohl (1997), especially in that both models 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. Finally, models of surprise have also been proposed in the fields of knowledge discovery and data mining (e.g. Suzuki & Kodratoff, 1998).

Macedo and Cardoso (e.g., Macedo & Cardoso, 2001)) developed a computational model of surprise that is an adaptation (although with several simplifications) of the models proposed by Meyer, Reisenzein and Schützwohl (1997) and by Ortony and Partridge (1987). In the present article, we elaborate and evaluate this model further by discussing different possible functions for the computation of surprise and by evaluating these functions in an empirical study.

The following section describes Macedo and Cardoso's surprise model in more detail, including an overview of its theoretical background models. Subsequently, we discuss several possible functions for computing the intensity of surprise. Finally, we describe an experimental test that was carried out to evaluate the accuracy of these surprise functions.

Surprise Model

As mentioned, the surprise model developed by Macedo and Cardoso (2001) is mainly based on Ortony and Partridge's (1987) proposals and on those of Meyer, Reisenzein and Schützwohl (1997). Therefore, we first give an overview of these background theories and then explain the computational model proposed by Macedo and Cardoso, by comparing it with these two models.

Background Models

Although Ortony and Partridge agree with Meyer, Reisenzein and Schützwohl and other authors that surprise is caused by events that are commonsensically called unexpected, they proposed that unexpectedness covers two cases. First, surprise results when prior expectations regarding an event are disconfirmed. Second, however, surprise can also be caused by events for which expectations were never computed. That is, according to Ortony and Partridge, there are situations in which one is surprised although one had no explicit expectations (either conscious or unconscious) regarding the surprising event. Ortony and Partridge also proposed that surprisingness is an important variable in artificial intelligence systems, particularly in attention and learning.

In more detail, Ortony and Partridge's model of surprise assumes a system (or agent) with an episodic and 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, they distinguish between practically deducible propositions 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. 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 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 preexisting representations or representations computed "after the fact". More precisely, surprise results in three situations (Table 1 presents 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*); 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, 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.

In their cognitive-psychoevolutionary model, Meyer, Reisenzein and Schützwohl also assume that surprise (considered by them as an emotion) is elicited by the appraisal of unexpectedness.

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*

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 (2001), 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; and (iv) possibly, immediate reactions to that event and/or updating or revision of the "old" schemas or beliefs.

Overview of the Computational Model of Surprise

Macedo and Cardoso (e.g., Macedo & Cardoso, 2001) developed a multi-agent environment in which, in addition to inanimate agents (objects such as buildings), there are two main kinds of animate, interacting agents: the "author-agents" or creators, whose main function is to create things (objects, events), and the "jury-agents" or explorers whose goal is to explore the environment by analyzing, studying and evaluating it. An agent can also show both of these activities (creation and exploration).

The computational model of surprise is integrated into the *motivations* module of the architecture of the artificial agents (see Figure 1). The other modules of this architecture are: *sensors/ perception*; *memory*; *goals/desires*; and *reasoning/decision-making*. This last module and the module *motivations* are provided with information from the world obtained through *sensors/perception*, as well as with information recorded in *memory*. The *reasoning/decision-making* module then computes the current state of the world. Afterwards, probability theory is applied to predict possible future states of the world for the available actions, and a utility function (which makes use of the intensity of the generated emotions) is applied to each of these world states. Finally, the action that maximizes the utility function is selected.

The computational model of surprise incorporated in this agent system is an adaptation (although with some simplifications) of the surprise model proposed by Meyer, Reisenzein and Schützwohl 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 Meyer, Reisenzein and Schützwohl model. Accordingly, in our model, we drew on the assumptions of Ortony and Partridge for the implementation of the appraisal of unexpectedness and the computation of the intensity of surprise, as well as for the selection of knowledge structures.

In Macedo and Cardoso's model, knowledge is exclusively of an episodic kind (for an example, see Figure 2), rather than being both semantic and episodic in nature (although this will be considered in future work), as in Ortony and Partridge's model. In this respect, the knowledge structure of our model also differs from the schema-theoretic framework of the Meyer, Reisenzein and Schützwohl model that also assumes both episodic and semantic knowledge. In our model, an input proposition (or new belief) is therefore always compared with episodic representations of objects or events (or their properties) (for instance an object with squared windows, rectangular door, etc.). Besides, the agent has in its episodic memory explicit representations of similar objects. 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 stored in episodic memory (see Figure 2). For instance, in the example shown in Figure 2, the proposition "houses have square facades" is immutable (since all the houses in memory have squared facades), whereas "houses have square windows" is a typical proposition with a probability (immutability) value of 0.50 (as implied by Ortony and Partridge's model, in our model immutability is a continuous variable).

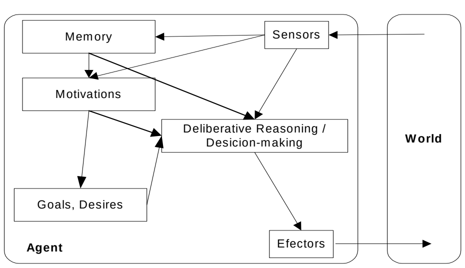

Figure 1: Architecture of an agent.

Field \ Case	C_1	C_2	C_3	C_4
Structure	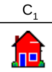		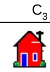	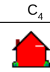
Function	House	House	Church	Hotel
Behavior	Static	Static	Static	Static
Abs. Freq.	50	40	5	5

Figure 2: Example of an episodic memory in the domain of buildings.

The usual activity of the agents consists of moving through the environment hoping to find interesting things (objects or events) that deserve to be investigated. We assume that this exploratory behavior is ultimately in the service of other (e.g., hedonic) motives, although this issue is not explicitly addressed in the present model. When one or more objects/events are perceived, the agent computes expectations for the missing information (e.g., "it is a house with 67% of probability", "it is a hotel with 45% of probability", etc.; note that the *function* of a building becomes available to the agent only when its position and that of the building are the same). On the basis of the available information (e.g., the visible *structure* of an object) and the computed expectations (e.g., predictions for the *function* of an object), the agent then determines the intensity of surprise that may be caused by the object/event (these computations, which correspond to the "appraisal of unexpectedness" in the Meyer, Reisenzein and Schützwohl model, are described in more detail below). Subsequently, the object/event with the maximum estimated surprise 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, Reisenzein and Schützwohl model. The previously estimated value of surprise may subsequently be updated on the basis of the additional information acquired about the object/event. The object/event is then 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, for instance, (Gärdenfors, 1988)).

The different surprise-eliciting situations distinguished by Ortony and Partridge are dealt with in our model in the following way. As said above, when an agent perceives an object, it first computes expectations (*deducible, active expectations*) for missing information (e.g., "it is a hotel with 45% of probability"). 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 (note that, in our model, belief conflicts may be partial as well as total). This is thus an example of the first source of surprise distinguished by Ortony and Partridge. In contrast, when an agent perceives an aspect or part of an object with particular properties (e.g., a building with a window of a circular shape) that were not actively predicted, it may still be able to infer that it expected something (e.g., a rectangular-shaped window with, 45% probability, a square-shaped window with 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”).

The Computation of Surprise Intensity

We now address the question of how the intensity of surprise should be computed in the model. In humans, this problem has already been successfully 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, 2000b) suggests that the intensity of felt surprise increases monotonically, and is closely correlated with, the degree of unexpectedness. On the basis of this evidence, we propose that the surprise “felt” by an agent elicited by an object/event X is proportional to the degree of unexpectedness of X (which in the model is based on the frequencies of objects/events present in the memory of the agent). According to probability theory, the degree of expecting an event X to occur is its subjective probability P(X). Accordingly, the improbability of X, denoted by 1-P(X), defines the degree of not expecting X, or for short its unexpectedness. The intensity of surprise elicited by X should therefore be an (at least weakly) monotonically increasing function of 1-P(X). As a first approach, this function (S1) could simply be taken to be the identity function, that is, the intensity of surprise could simply be equated with the degree of unexpectedness:

$$S1(Agt, X) = 1 - P(X)$$

However, on second thought, S1 does not seem to faithfully capture the relation between unexpectedness and surprise. For example, consider a political election with three candidates A, B and C, where the probability of being elected is P(A) = P(B) = P(C) = 0.333. In this case, one would not be surprised if either A, B or C is elected. Therefore, in this situation at least, S1 fails.

To arrive at a more adequate surprise function, consider the case where there are only two mutually exclusive and exhaustive alternative events, X and Y (i.e., not X). Here, intuition suggests that X is not surprising as long as P(X) ≥ 0.5, whereas X is surprising for P(X) < 0.5, and increasingly more so the more P(X) approaches 0. This intuition is captured by the following surprise function (S2):

$$S2(Agt, X) = \begin{cases} 1 - P(X) \Leftarrow P(X) < 0.5 \\ 0 \Leftarrow P(X) \geq 0.5 \end{cases}$$

To deal with sets of more than two mutually exclusive events, S2 could be generalized as follows (S3, where *n* denotes the number of events in the set):

$$S3(Agt, X) = \begin{cases} 1 - P(X) \Leftarrow P(X) < \frac{1}{n} \\ 0 \Leftarrow P(X) \geq \frac{1}{n} \end{cases}$$

However, it may be more adequate to set the upper limit of surprise not to 1, but to $\frac{1}{n}$ (see S4):

$$S4(Agt, X) = \begin{cases} \frac{1}{n} - P(X) \Leftarrow P(X) < \frac{1}{n} \\ 0 \Leftarrow P(X) \geq \frac{1}{n} \end{cases}$$

Yet another possible surprise function, suggested by further reflection on the above election example, is the following (S5):

$$S5(Agt, X) = P(Y) - P(X)$$

In this formula, Y is the event with the *highest* probability of a set of mutually exclusive events. S5 implies that, within each set of mutually exclusive events, there is always one (Y) whose occurrence is entirely unsurprising, namely the event with the maximum probability in the set (P(Y)). For the other events X in the set, the surprise intensity caused by their occurrence is the difference between P(Y) and their probability P(X). This difference can be interpreted as the amount by which P(X) has to be increased for X to become unsurprising. For instance, in the election example considered earlier, where P(A) = P(B)= P(C) = 0.333, S5 correctly predicts that one would not be surprised if either A, B or C is elected. By contrast, if P(A) = 0.55, P(B) = 0.40 and P(C) = 0.05, S5 predicts that the surprise caused by B is 0.15 and for C is 0.50, whereas for A it is 0. S5 also implies that maximum surprise, that is, S(X) = 1, occurs only if P(Y) = 1 and hence, by implication, P(X) = 0. (In the Ortony and Partridge model, this corresponds to situations [1], [2], [5] and [6], where the disconfirmed event Y is immutable, i.e., its probability is 1). Therefore, S5 seems to correctly describe surprise in the election example. Confirming this impression, S5 also acknowledges the intuition behind S2: if there are only two alternative events X and Y (= not X), S5 predicts, like S2, that X should be unsurprising for P(X) ≥ 0.5, for in this case X is also the event with the highest probability in the set. By contrast, for P(X) < 0.5, S5 predicts that X should be surprising and increasingly so the more P(X) approaches 0, with maximum possible surprise (S(X) = 1) being experienced for P(X) = 0.

Yet another possible surprise function (S6) is suggested by Information Theory (Shannon, 1948):

$$S6(Agt, X) = \log_2 \frac{1}{P(X)}$$

According to S6, surprise about X is 0 when P(X) = 1 and increases monotonically with decreasing P(X). In these respects, then, S6 is similar to S1. However, in contrast to S1, S6 is a nonlinear function of P(X), and it is not normalized. For instance, for P(X) = 0.3, *S6(X)* = 1.7 (bits), for P(X) = 0.01, *S6(X)* = 6.6, and for P(X) = 0.001, *S6(X)* = 9.9. In fact, there is no upper limit of S(X): for P(X)=0, *S6(X)* = $+\infty$. To overcome this problem, we propose the following normalized function S7 (stipulating the upper limit to be 10):

$$S7(Agt, X) = \frac{\log_2 \frac{1}{P(X)}}{10}$$

Finally, yet another surprise function (S8), a nonlinear modification of S5, is suggested by the results of the experiment, reported below, performed with humans in the domain of elections and sport games:

$$S8(Agt, X) = \log_2 (1 + P(Y) - P(X))$$

This function retains the essential features of S5: when X is the most expected event (X = Y), then S8(X) = 0; when X is different from Y, S8(X) > 0 and increases monotonically with the difference between P(Y) and P(X); and S8(X) is maximal (= 1) if P(Y) = 1 and P(X) = 0. In addition, however, S8 also captures the nonlinearity of the surprise function suggested by the experiments with humans reported below.

Experiment

To test the validity of the proposed surprise functions, we conducted an experiment that involved two steps. In step 1, we collected ratings of probability and surprise intensity from humans in two domains, political elections and sports games. In step 2, artificial agents that implemented the different surprise functions were provided with the probability judgments obtained from the humans and, on this basis, computed surprise intensity values. These predicted surprise values were then compared with the actual surprise ratings provided by the human participants.

Step 1 was conducted with ten participants (mean age, 29 years). They were presented with 20 brief scenarios, 10 of which described political elections with 2-4 candidates (see Figure 3), whereas the other 10 scenarios described sports games with 2-4 teams or players (see (Reisenzein, 2000a) for a conceptually similar experiment using knowledge questions). Political elections and sports games were chosen because we thought that these domains are familiar to most people and that the participants would have no problems to state their probabilities and their surprise about outcomes. In addition, in contrast to the domain of buildings used in a previous study reported in (Macedo & Cardoso, 2001), elections and sport games allow for an easier matching of the knowledge of artificial agents with that of humans. Part of the scenarios did not include information about the actual outcome of the election or game, whereas the remaining scenarios included this information. For scenarios without outcome information, the participants were asked to first state their expectations for all possible outcomes and to rate their probability on a 1-100 scale. Subsequently, they were informed about the outcome of the election or game and rated their surprise about the outcome first on a qualitative intensity scale and then again on a quantitative intensity scale within the chosen qualitative level. By contrast, for the scenarios that included outcome information, participants first rated the intensity of surprise about the outcome and subsequently their (passive) expectations regarding the outcome. An example of a scenario is shown in Figure 3.

> Given the following prognosis for the election of candidate A, B and C for a political position:
>
> Victory of A=45%; Victory of B=45%; Victory of C=10%
>
> a) What are your personal expectations regarding the victory of candidates A, B and C?
> b) Assume that candidate A won the election and rate the intensity of surprise that you would feel.

Figure 3: Example of a test item.

In step 2 of the study, the probability ratings obtained from each participant in step 1 were delivered to eight artificial agents, each of which implemented one of the eight surprise functions S1-S8 described earlier. Using these functions, the agents computed surprise intensity values from the probabilities. These predicted surprise values were then compared with the surprise ratings of the humans obtained in step 1.

The data obtained in the first step of the experiment suggested two qualitative conclusions. First, the occurrence of the most expected event of the set of mutually exclusive and exhaustive events did not elicit surprise in humans. For example, when the expectations for the election of three political candidates A, B and C were P(A) = 0.55, P(B) = 0.40, and P(C) = 0.05, the participants felt no surprise about the election of candidate A. This was also true when two or more candidates had equal maximal probabilities. For example, when P(A) = 0.40, P(B) = 0.40 and P(C) = 0.20, participants were not surprised when either A or B was elected. Second, beyond the point of zero surprise, the surprise function appeared to be nonlinear. For example, relatively high surprise was indicated when candidate C won the elections in both of the above situations, although it was still higher for P(C) = 0.05 than for P(C) = 0.20.

To compare the surprise values generated by the artificial agents and the surprise ratings provided by the human judges, the following fit indices were used: the root mean squared difference, the mean absolute difference, and the Pearson correlation. The results of these comparisons are shown in Table 2, separately for the 10 participants (H1, …, H10) and for six of the eight artificial agents (A1,…,A8) (the surprise functions S6 and S7 were not included because they have a different range than the human ratings and therefore computation of the absolute and squared differences is not meaningful). It can be seen from Table 2 that, regardless of which fit index is used, agent A8 (which implemented surprise function S8) was the one with the best fit to the human ratings: it had on average, the lowest root mean squared differences (*Ms*= 0.10), the lowest absolute differences (*Md*= 0.06), and the highest correlation to these ratings (*Mr*= 0.98). A8 was closely followed by A5 (*Ms* = 0.21; *Md* = 0.08; *Mr* = 0.97), whereas agents A1 and A2 had the comparatively worst fit values (for instance, A1 had *Ms* = 0.35; *Md* = 0.26; *Mr* = 0.81). A main reason for the bad performance of A1 was apparently that it failed in the case of the occurrence of the most expected event of the set: A1 still predicts a positive surprise value (1-P(X)) for this case, whereas humans do not feel surprised by the occurrence of this event. However, in other situations, A1 performed well.

Table 2: Statistical comparison of the surprise values computed by the artificial agents and those provided by the humans (*s* = root mean squared difference, *d* = mean absolute difference, and *r* = Pearson correlation).

		H1	*H2*	*H3*	*H4*	*H5*	*H6*	*H7*	*H8*	*H9*	*H10*	*M*
A1	*s*	.35	.36	.34	.35	.35	.34	.35	.36	.35	.36	**.35**
	d	.25	.26	.25	.25	.26	.24	.27	.27	.26	.27	**.26**
	r	.82	.80	.82	.82	.80	.82	.81	.80	.82	.82	**.81**
A2	*s*	.30	.33	.29	.32	.32	.30	.33	.32	.31	.31	**.31**
	d	.18	.21	.16	.20	.21	.18	.22	.19	.19	.19	**.19**
	r	.82	.79	.82	.81	.79	.83	.80	.80	.81	.81	**.81**
A3	*s*	.22	.30	.24	.21	.30	.22	.18	.19	.19	.16	**.22**
	d	.07	.15	.09	.07	.17	.09	.09	.09	.08	.08	**.10**
	r	.95	.85	.89	.94	.81	.92	.93	.92	.92	.94	**.91**
A4	*s*	.43	.41	.45	.43	.43	.43	.44	.46	.46	.45	**.44**
	d	.29	.28	.30	.29	.29	.28	.28	.28	.29	.27	**.28**
	r	.93	.92	.88	.96	.90	.95	.91	.91	.93	.94	**.92**
A5	*s*	.22	.16	.19	.16	.23	.20	.21	.24	.24	.24	**.21**
	d	.07	.06	.11	.06	.09	.05	.08	.10	.09	.09	**.08**
	r	.97	.98	.96	.98	.95	.99	.97	.96	.96	.96	**.97**
A8	*s*	.09	.07	.13	.08	.12	.06	.11	.13	.12	.12	**.10**
	d	.05	.05	.09	.05	.08	.04	.06	.08	.07	.07	**.06**
	r	.98	.99	.98	.99	.97	.99	.98	.07	.07	.97	**.98**

Conclusions

The empirical study of the surprise functions suggests S8(X) = $\log_2(1+P(Y)-P(X))$ as the most appropriate surprise function for the domains of political elections and sport games, although S5 (the linear counterpart of S8) is a very close contender. However, before more definitive conclusions can be drawn, additional tests need to be performed in other domains, as well as with yet other possible surprise functions (e.g., Shackle, 1969).

Acknowledgments

We would like to thank A. Ortony for his helpful comments. The PhD of Luís Macedo is financially supported by PRODEP III.

References

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

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

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

Izard, C. (1991). The Psychology of Emotions. NY: 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.

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. (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. 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.

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

Shackle, G. (1969). Decision, Order and Time in Human Affairs (2 ed.). Cambridge, UK: Cambridge University Press.

Shannon, C. (1948). A mathematical theory of communication. Bell System Technical Journal, 27, 379-423;623-656.

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.

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