A Computational Model of Selective Attention

Luis Macedo (macedo@dei.uc.pt)
University of Coimbra
Centre for Informatics and Systems of the University of Coimbra, Coimbra, Portugal

Abstract

Inspired on natural selective attention studies, we propose a computational model of selective attention that relies on the assumption that uncertain, surprising and motive congruent/incongruent information demands attention from an intelligent agent. This computational model has been integrated into the architecture of a Belief-Desire-Intention artificial agent so that this can autonomously select relevant, interesting information of the (external or internal) environment while ignoring other less relevant information. The advantage is that the agent can communicate only that interesting, selective information to its processing resources (focus of the senses, decision-making, etc.) or to its human owner's processing resources so that these resources can be allocated more effectively. We illustrate this role of the artificial, selective attention mechanism in the time-critical, risky situation, of driving a vehicle, by showing that it prevents both the personal traffic assistant agent's and its human owner's decision-making resources of receiving unnecessary traffic information.

Keywords: Selective attention; Interest; Value of information, Surprise; Uncertainty; Information overload; Resource-bounded agents; Personal agents.

Introduction

In many ways, the advent of information technology is a primary reason for the abundance of information with which humans are inundated, due to its ability to produce more information more quickly and to disseminate this information to a wider audience than ever before. Contrary to what in general could be expected, a lot of recent studies confirmed what Alvin Toffler (1970) predicted a few decades ago: the overabundance of information instead of being beneficial is a huge problem having many negative implications not only in personal life but also in organizations, business, and in general in the world economy. Research proves that the brain simply does not deal very well with this multitasking process: there is a waste of time as the brain switches from one task to another and back again (Klingberg, 2008). This explains why decision quality and the rate of performing tasks degrades with increases in the amount of information being considered.

A fundamental strategy for dealing with this problem of information overload (O'Connell, 2008) should include making devices that incorporate themselves selective attention agents in order to decrease the amount of information considered in their own reasoning/decision-making processes or decrease the amount of information provided by them to humans, preventing these from a number of interruptions.

But how to model selective attention in artificial agents? The problem starts at the human level. Although selective attention has been thoroughly researched over the last 100 years in psychology and more recently in neuroscience (e.g., Kahneman, 1973; Wright & Ward, 2008), at present there is no general theory of selective attention. Instead there are specific theories for specific tasks such as orienting, visual search, filtering, multiple action monitoring (dual task), and multiple object tracking.

In spite of this, a number of models of selective attention has been proposed in Cognitive Science (e.g., Horvitz, Jacobs, & Hovel, 1999). Particularly related with these models is the issue of measuring the value of information. A considerable amount of literature has been published on these measures, especially from the fields of active learning and experimental design. Most of those measures rely on assessing the utility or the informativeness of information (e.g., Horvitz & Barry, 1995; MacKay, 1992; Lindley, 1955; Settles, 2008) However, little attention has been given to the surprising and motive congruence value of information, giving the beliefs and desires of an agent.

Opposed to other approaches (e.g., Itti & Baldi, 2006; Peters, 1998; Schmidhuber, 2006; Oudeyer, Kaplan, & Hafner, 2007) relying on low-level, raw information, Macedo, Cardoso, and Reisenzein (2001; 2004), and Lorini and Castelfranchi (2007) proposed, independently, computational models of surprise that are based on the mechanism that compares newly acquired beliefs to preexisting beliefs. Both models of artificial surprise were influenced by psychological theories of surprise (e.g., Meyer, Reisenzein, & Schützwohl, 1997), and both seek to capture essential aspects of human surprise (see Macedo, Cardoso, Reisenzein, Lorini, & Castelfranchi, 2009, for a comparison). In agreement with most theories of human surprise, both models of artificial surprise conceptualize surprise as a fundamentally expectation- or belief-based cognitive phenomenon, that is, as a reaction to the disconfirmation of expectations or, more generally, beliefs. Furthermore, in both models, beliefs are understood as propositional attitudes (e.g., Searle, 1983), and a quantitative belief concept (subjective probability) is used. Both artificial surprise models draw a distinction between two main kinds of expectations or beliefs whose disconfirmation causes surprise (see also Ortony & Partridge, 1987): Active versus passive expectations. Although Macedo and Cardoso initially used the same surprise intensity function, according to which the intensity of surprise about an event is proportional to its unexpectedness, Macedo, Reisenzein and Cardoso subsequently opted for a "contrast model" of surprise intensity. This model assumes that the intensity of surprise about an event reflects its probability difference to the contextually most expected event (see also Teigen & Keren, 2003).

In this paper we describe an artificial selective attention mechanism that may be used by artificial agents so that only cognitively and affectively, interesting/relevant information is selected and forwarded to reasoning/decision-making units. Our approach relies on the psychological and neuroscience studies about selective attention which defend that variables such as unexpectedness, unpredictability, surprise, uncertainty, and motive congruence demand attention (e.g., Berlyne, 1960; Kahneman, 1973; Ortony & Partridge, 1987). One of the features of the selective attention mechanism is that it should work in the absence of a model of decision-making of the artificial agent, or of its designer, owner or user for whom the artificial agent might act on his/her behalf.

The next section describes the computational model of selective attention, focusing on how the multidimensional value of information is computed, which will be illustrated with an example in Section 3 . Section 4 examines the performance of the selective attention mechanism as well as its role on the decrease of unnecessary information while not affecting significantly an agent decision-making performance. Finally, in Section 5 we present conclusions.

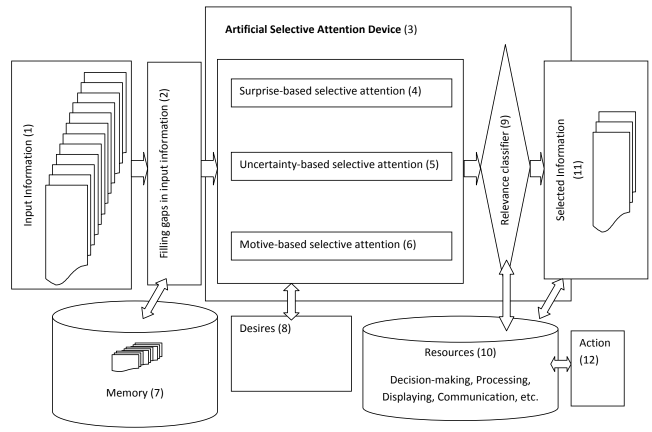

Figure 1: Architecture of an artificial selective attention agent.

A Computational Model for Forms of Selective Attention

Selective attention may be defined as the cognitive process of selective allocation of processing resources (focus of the senses, etc.) on relevant, important or interesting information of the (external or internal) environment while ignoring other less relevant information. The issue is how to measure the value of information. What makes something interesting? In cognitive science, attentional focus is linked with expectation generation and failure, i.e., with surprise (Ortony & Partridge, 1987). Therefore, it is reasonable to consider that any model of selective attention should rely on a cognitive model of surprise. However, surprise is not enough. Happiness/pleasantness, which according to cognitive theories of emotion and specifically to belief-desire theories of emotion (Reisenzein, 2008) is directly related to congruence between new information and the human agent's motives/desires, may also play also a fundamental role on attention. For this reason, the system must also incorporate a measure of the expected satiation of the desires.

In order to accomplish all those requirements, we developed an architecture for a personalized, artificial selective attention agent (see Figure 1). We assume: (i) this agent interacts with the external world receiving from it information through the senses and outputs actions through their effectors; (ii) the world is described by a large amount of statistical experiments; (iii) the agent is a BDI agent (Rao & Georgeff, 1995), exhibiting a prediction model (model for generating expectations, i.e., beliefs about the environment), a desire strength prediction model (a model for generating desire strengths for all the outcomes of the statistical experiments of the world that are know given the desires of the agent – profile of the agent which include basic desires), as well as the intentions (these define the profile of the agent); (iv) the agent contains other resources for the purpose of reasoning and decision-making.

While the belief strengths are inferred from data using a frequentist approach and updated as new information is acquired, the desirability of the outcomes is previously set up although they depend on the intention of the agent, suffering changes whenever the agent is committed with a new intention.

The first of the modules of the architecture (module 1 in Figure 1) is concerned with getting the input information. The second is the computation of the current world state. This is performed by generating expectations or assumptions (module 2), based on the knowledge stored in memory, for the gaps of the environment information provided by the sensors (module 1). We assume that each piece of information resulting from this process, before it is processed by other cognitive skills, goes through several sub-selective attention devices, each one evaluating information according to a certain dimension such as surprise (module 4), uncertainty (module 5), and motive-congruence/incongruence – happiness (module 6). For this task the selective attention mechanism takes into account some knowledge container (memory — preexisting information (module 7)), and the intentions and desires (motives — module 8). There is a decision-making module (module 9) that takes into account the values computed by those sub-selective attention modules and decides if a piece of information is relevant/interesting or not. Then, this module of decision-making selects the higher relevant pieces of information so that other resources (reasoning, decision-making, displaying, communication resources, etc.) (module 10) can be allocated to deal with them.

The representation of the memory contents (beliefs) relies on semantic features or attributes much like in semantic networks (Russell & Norvig, 2010) or schemas (Rumelhardt & Ortony, 1977). Each attribute, $attr_i$, viewed by us as a statistical experiment, is described by a probabilistic distribution, i.e., a set $A_i = \{< value_j, prob_j, desireStrength_j >: j = 1,2,\ldots,n\}$, where n is the number of possible values of the attribute, $P(attr_i = value_j) = prob_j$, and $desireStrength_j$

is the desirability of $attr_i = value_j$ (for a related work see (Reisenzein, 2008)).

The next sub-sections describe each one of the dimensions for evaluating information, namely surprise, uncertainty, and motive congruence/incongruence. While the dimensions of surprise and uncertainty are related to the value of information to the belief store of the agent, the dimension of motive congruence/incongruence is related to the value of information to the goals/desires of the agent (these dimensions are related to the concepts of cognitive and affective feelings of (Clore, 1992) and belief-belief and belief-desire comparators of (Reisenzein, 2008)).

Surprise Value of Information

We adopted the computational model of surprise of (Macedo & Cardoso, 2001; Macedo et al., 2004) which is formally defined in Definition 1 (for related models see Macedo et al., 2009). Macedo, Cardoso and Reisenzein computational model of surprise suggests that the intensity of surprise about an event E_g, from a set of mutually exclusive events $E_1, E_2, \dots, E_m$, is a nonlinear function of the difference, or contrast, between its probability and the probability of the highest expected event E_h in the set of mutually exclusive events $E_1, E_2, \dots, E_m$.

Definition 1 *Let (Ω, A, P) be a probability space where Ω is the sample space (i.e., the set of possible outcomes of the experiment), $A = A_1, A_2, .., A_n$ is a σ-field of subsets of Ω (also called the event space, i.e., all the possible events), and P is a probability measure which assigns a real number $P(F)$ to every member F of the σ-field A. Let $E = \{E_1, E_2, \dots, E_m\}$, $E_i \in A$, be a set of mutually exclusive events in that probability space with probabilities $P(E_i) >= 0$, such that $\sum_{i=1}^{m} P(E_i) = 1$. Let E_h be the highest expected event from E. The intensity of surprise about an event E_g from E is given by:*

$$S(E_g) = \log(1 + P(E_h) - P(E_g)) \qquad (1)$$

The probability difference between $P(E_h)$ and $P(E_g)$ can be interpreted as the amount by which the probability of E_g would have to be increased for E_g to become unsurprising.

Proposition 1 *In each set of mutually exclusive events, there is always at least one event whose occurrence is unsurprising, namely, E_h.*

Uncertainty-based Value of Information

Information is a decrease in uncertainty which, according to information theory, is measured by entropy (Shannon, 1948). When new information is acquired its amount may be measured by the difference between the prior uncertainty and the posterior uncertainty.

Definition 2 *Let (Ω, A, P_{prior}) be a probability space where Ω is the sample space (i.e., the set of possible outcomes of the experiment), $A = A_1, A_2, .., A_m$ is a σ-field of subsets of Ω (also called the event space, i.e., all the possible events), and P_{prior} is a probability measure which assigns a real number $P_{prior}(F)$ to every member F of the σ-field A. Let $E = \{E_1, E_2, \dots, E_m\}$, $E_i \in A$, be a set of mutually exclusive events in that probability space with probabilities $P_{prior}(E_i) >= 0$, such that $\sum_{i=1}^{m} P_{prior}(E_i) = 1$. Let P_{post} be the posterior probability measure, after some data is acquired, which assigns a real number $P_{post}(F)$ to every member F of the σ-field A such that it assigns $P_{post}(E_i) >= 0$ with $\sum_{i=1}^{m} P_{post}(E_i) = 1$. According to information theory, the information gain of an agent after some data is acquired, $IG(E)$, is given by the decrease in uncertainty:*

$$\begin{aligned} IG(E) &= H_{prior}(E) - H_{post}(E) \\ &= -\sum_{i=1}^{m} P_{prior}(E_i) \times \log(P_{prior}(E_i) - \\ &\quad (-\sum_{i=1}^{m} P_{post}(E_i) \times \log(P_{post}(E_i)) \qquad (2) \end{aligned}$$

$H_{post} = 0$ if and only if all the $P_{post}(E_i)$ but one are zero, this one having the value unity. Thus only when we are certain of the outcome does H_{post} vanish, otherwise it is positive.

IG is not normalized. In order to normalize it we must divide it by $\log(m)$ since it can be proved that $IG \leq \log(m)$:

$$IG(E) = \frac{H_{prior}(E) - H_{post}(E)}{\log(m)} \qquad (3)$$

Motive Congruence/Incongruence-based Value of Information

While the measure of surprise takes into account beliefs that can be confirmed or not, the pleasantness function that we describe in this subsection takes as input desires that, contrary to beliefs, can be satisfied or frustrated. Following the belief-desire theory of emotion (Reisenzein, 2008), we assume that an agent feels happiness if it desires a state of affairs (a proposition) and firmly beliefs that that state of affairs obtains. The intensity of happiness about an event is a monotonically increasing function of the degree of desire of that event as formally defined in Definition 4.

Definition 3 *Let (Ω, A) be a measurable space where Ω is the sample space (i.e., the set of possible outcomes of the experiment) and $A = A_1, A_2, .., A_m$ a σ-field of subsets of Ω (also called the event space, i.e., all the possible events). We define the measure of desirability of an event on (Ω, A) as $D : A \to [-1, 1]$, i.e., as a signed measure which assigns a real number $-1 \leq D(F) \leq 1$ to every member F of the σ-field A based on the profile of the agent, so that the following properties are satisfied:*

- $D(\emptyset) = 0$, $-|\Omega| \geq D(\Omega) \leqslant |\Omega|$
- *if $A_1, A_2, \dots$ is a collection of disjoint members of A, in that $A_i \cap A_j = \emptyset$ for all $i \neq j$, then*

$$D(\bigcup_{i=0}^{\infty} A_i) = \sum_{i=0}^{\infty} D(A_i) \quad (4)$$

The triple (Ω, A, D) is called the desirability space.

Definition 4 *Let (Ω, A, P) and (Ω, A, D) be the probability and the desirability spaces described, respectively, in Definition 1 and Definition 3. Let $E = \{E_1, E_2, \ldots, E_m\}$, $E_i \in A$, be a set of mutually exclusive events in that probability space with probabilities $P(E_i) >= 0$, $\sum_{i=1}^{m} P(E_i) = 1$. If $P(E_g) = 1$, the intensity of happiness, i.e., motive congruence, about an event E_g from E is given by:*

$$MC(E_g) = D(E_g) \quad (5)$$

The Principle of Selective Attention

Having defined the motive, the uncertainty-based, and surprise-based selective attention modules, we are now in a position to formulate, in a restricted sense (without the inclusion of other information measures such as complexity), the principle that a resource-bounded rational agent should follow in order to avoid an overabundance of information and interruptions in the absence of a model for decision-making. Note that if this model is known, the problem is reduced to the classical computation of the value of information that has been extensively studied (e.g., Horvitz & Barry, 1995; Russell & Norvig, 2010).

Definition 5 *A resource-bounded rational agent should focus its attention only on the relevant and interesting information, i.e., on information that is congruent or incongruent to its motives/desires, and that is cognitively relevant because it is surprising or because it decreases uncertainty.*

We may define real numbers α, β, and γ as levels above which the absolute values of motive congruency, surprise, and information gain (decrease of uncertainty), respectively, should be so that the information can be considered valuable or interesting. These are what we called the triggering levels of alert of the selective attention mechanism. Note that, making one of those parameters null is equivalent to removing the contribution of the corresponding component from the selective attention mechanism.

Practical Application

Advanced Travel Information Systems (ATIS) are designed to assist travelers in making pre-trip and en-route travel decisions by providing them pre-trip and en-route information. However, while these information systems can undoubtedly help humans perform better in these complex traveling scenarios, they might provide an unhandled amount of information to humans that may compromise their performance.

It is contended that while many traveler information systems are innovative and make use of cutting edge technologies, they lack real machine intelligence and therefore may be limited in their ability to service the traveling public over the long-run. On the one hand, a wave of technological developments, in particular the increasing deployment of GIS and, on the other hand, the introduction and rapid market penetration of mobile devices such as cell phones boosted the development of ATIS towards what has been termed Intelligent Traveler Information Systems (ITIS) (Adler & Blue, 1998), in which artificial intelligence techniques are drawn upon to create systems capable of providing travelers with more personalized planning assistance. This is the goal of integrating selective agents in personal devices that receive information from the ATIS to act as personal assistant selective attention agents in order to avoid unnecessary interruptions to their users by enabling that only interesting information (i.e., with a value above a threshold defined by the user) is provided to them.

We are developing an ITIS that receives information about the traffic conditions and sends it to the mobile devices of the travelers. All that collected information is stored in the knowledge base/memory of the ITIS. There is a personal selective attention agent for each registered traveler. Each one of these personal agents has information about the expectations of its owner based on their travel history.

Let us illustrate how the value of information is computed by the selective attention mechanism. Suppose that a traveller's navigation system provided the pre-route path containing a road A for an agent (a driver) based on its profile (e.g., preference for shortest routes). Suppose the agent has the following expectations for the traffic conditions of road A, for a certain period/time of the day for a certain day of the week: 60% of probability of "good traffic conditions" (event E_1), 30% of probability of "moderate traffic conditions" (event E_2), and 10% of probability of "bad traffic conditions" (event E_3). Suppose the desire strengths of these events are 1, -0.5, and -1, respectively. Given that the agent plans to go trough that route, suppose its module for generating/managing desires assigns a null desire strength for the other routes as it does not care about the traffic conditions of the other roads that are not part of its planned route. What is the relevance of becoming aware that the current traffic conditions of road A are good (event E_1)? Considering solely the motive-based component, the outcomes (events E_1, E_2, and E_3) elicits happiness (motive congruence) with intensity 1, -0.5 and -1, respectively. E_1 is congruent/consistent with the goals of the agent, while E_2 and E_3 are incongruent with the goals of the agent.

According to Equation 1, the surprise value of E_1, E_2, and E_3 are, respectively, 0, 0.38, and 0.58. Illustrating for the case of E_3:

$$\begin{aligned} Surprise(E_3) &= \log(1 + P(E_1) - P(E_3)) \\ &= \log(1 + 0.6 - 0.1) = 0.58 \end{aligned} \quad (6)$$

According to Equation 3, the normalized information gain value of E_1, E_2, or E_3 is:

$$IG(E) = \frac{H_{prior}(E) - H_{post}(E)}{\log(m)} = \frac{H_{prior}(E) - 0}{\log(3)}$$
$$= \frac{-\sum_{i=1}^{3} P_{prior}(E_i) \times \log(P_{prior}(E_i))}{\log(3)} = 0.82 \quad (7)$$

Assume the Principle of Selective Attention described above, with parameters $\alpha = 0.3$, $\beta = 0.5$, and $\gamma = 0.6$. Are all these events interesting? Considering the motive-based component all those events are interesting. However, from the perspective of the surprise-based selective attention component, the answer is "no" to the question related with the events E_1 and E_2 in that their surprise values, 0 and 0.38, respectively, are below β. With respect to E_3 the answer is "yes" given that its surprise value is 0.58. Taking the uncertainty-based component into account, the answer is "yes" for all the events because their occurrence gives a normalized information gain of 0.82 which is above γ.

Discussion and Conclusions

We presented a computational model for selective attention based on cognitive and affective feelings. We found evidence indicating that the mechanism contributes for decreasing the amount of unnecessary information while maintaining acceptable the performance of the owner (a human).

The advantages of reasoning correctly with less information include spending less time in processing information which is important in time-critical, high-risk situations. Besides, agents equipped with a selective attention filter can be successful personal assistants of humans, integrated for instance in mobile devices, so that their human users are prevented from unnecessary interruptions. This may be of high value in critical situations such as driving a car in that, as reported by (Horvitz & Barry, 1995), numerous cognitive studies have provided evidence of the problems in information processing exhibited by humans when dealing with large amounts of information such as that the speed at which humans perform tasks drops as the quantity of information being considered increases, and that the rate of performing tasks can be increased by filtering irrelevant information. In this particular case of transportation systems, the ultimate advantage may be less vehicle accidents and less deaths, while in organizations the advantage may be an improvement in their workers productivity and therefore less costs.

An hypothesis that might be risen is that taking other sub-selective attention modules such as those based on other cognitive or affective feelings (Clore, 1992) (e.g., familiarity, complexity) into account improves the performance of the mechanism. More experiments should be done with this aim.

Acknowledgments

Work funded by Fundação para a Ciência e Tecnologia — Project PTDC/EIA-EIA/108675/2008, and by FEDER through Programa Operacional Factores de Competitividade do QREN — COMPETE:FCOMP-01-0124-FEDER-010146.

References

Adler, J., & Blue, V. (1998). Toward the design of intelligent traveler information systems. *Transportation Research Journal Part C*, *6/3*, 157-172.

Berlyne, D. (1960). *Conflict, arousal and curiosity*. New York: McGraw-Hill.

Clore, G. (1992). Cognitive phenomenology: Feelings and the construction of judgment. In L. Martin & A. Tesser (Eds.), *The construction of social judgments* (p. 133-163). Hillsdale, NJ: Lawrence Erlbaum Associates.

Horvitz, E., & Barry, M. (1995). Display of information for time-critical decision-making. In *Proceedings of the eleventh conference on uncertainty in artificial intelligence* (p. 296-305). Morgan Kaufmann.

Horvitz, E., Jacobs, A., & Hovel, D. (1999). Attention-sensitive alerting. In *Proceedings of the fifteenth conference on uncertainty and artificial intelligence* (p. 305-313). Morgan Kaufmann.

Itti, L., & Baldi, P. (2006). Bayesian surprise attracts human attention. *Advances in Neural Information Processing Systems (NIPS 2005)*, *19*, 1-8.

Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.

Klingberg, T. (2008). *The overflowing brain: information overload and the limits of working memory*. Oxford University Press.

Lindley, D. (1955). On a measure of information provided by an experiment. *The Annals of Mathematical Statistics*, 986-1005.

Lorini, E., & Castelfranchi, C. (2007). The cognitive structure of surprise: looking for basic principles. *Topoi: An International Review of Philosophy*, *26*(1), 133-149.

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* (p. 588-593). Edinburgh, Scotland, UK: Erlbaum.

Macedo, L., Cardoso, A., Reisenzein, R., Lorini, E., & Castelfranchi, C. (2009). Artificial surprise. In J. Vallverd & D. Casacuberta (Eds.), *Handbook of research on synthetic emotions and sociable robotics: New applications in affective computing and artificial intelligence* (p. 267-291). Hershey: USA: IGI Global.

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* (pp. 873–878). Chicago, Illinois, USA: Lawrence Erlbaum Associates, Inc.

MacKay, D. (1992). Information-based objective functions for active data selection. *Neural Computation*, *4*(4), 590-604.

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

O'Connell, N. (2008). Interruption overload. *Strategic Direction*, *24*(10), 3-5.

Ortony, A., & Partridge, D. (1987). Surprisingness and expectation failure: what's the difference? In *Proceedings of the 10th international joint conference on artificial intelligence* (p. 106-8). Milan, Italy: Morgan Kaufmann.

Oudeyer, P., Kaplan, F., & Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, *11*(2), 265–286.

Peters, M. (1998). Towards artificial forms of intelligence, creativity, and surprise. In *Proceedings of the twentieth annual conference of the cognitive science society* (p. 836-841). Madison, Wisconsin, USA: Erlbaum.

Rao, A., & Georgeff, M. (1995). Bdi agents: from theory to practice. In *Proceedings of the first international conference on multiagent systems* (p. 312-319). San Francisco, CA, USA: MIT Press.

Reisenzein, R. (2008). Emotions as metarepresentational states of mind: Naturalizing the belief-desire theory of emotion. *Cognitive Systems Research*, *9*.

Rumelhardt, D., & Ortony, A. (1977). The representation of knowledge in memory. In R. Anderson, R. Spiro, & W. Montague (Eds.), *Schooling and the acquisition of knowledge* (p. 99-135). Hillsdale, NJ: Lawrence Erlbaum Associates.

Russell, S., & Norvig, P. (2010). *Artificial intelligence - a modern approach* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.

Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, *18*, 173-187.

Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.

Settles, B. (2008). *Curious machines: Active learning with structured instances*. Unpublished doctoral dissertation, University of Wisconsin, Madison, USA.

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

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

Toffler, A. (1970). *Future shock*. Bantam Book.

Wright, R. D., & Ward, L. M. (2008). *Orienting of attention*. Oxford, UK: Oxford University Press.