Curiosity: From Psychology to Computation

Qiong Wu and Chunyan Miao, Nanyang Technological University

Literature in psychology has shown that curiosity is the intrinsic motivation for exploration, learning and creativity. Various forms of computational curiosity have been developed to provide artificial beings with desirable functions such as detecting and adapting to novel inputs, making decisions related to aesthetics, and achieving pedagogical purposes. This paper reviews existing models of computational curiosity in light of psychological theories which are beneficial to building models of human cognition and designing human-like agents. We first study theories in psychology to shed light on the underpinnings of human curiosity, where a two-step process is proposed to serve as a general model for analyzing curiosity. Subsequently, existing models of computational curiosity are reviewed under the proposed framework. We conclude the review by identifying four key research issues in computational curiosity and ten important research areas where computational curiosity could bring significant impacts.

Categories and Subject Descriptors: I.2.0 [**Artificial Intelligence**]: General-*Cognitive simulation*; I.2.6 [**Artificial Intelligence**]: Learning; I.2.9 [**Artificial Intelligence**]: Robotics; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence-*Intelligent agents*

General Terms: Theory, Model, Application

Additional Key Words and Phrases: Human Curiosity, Computational Curiosity, Curious Agents, Curious Robots, Curious Learning

ACM Reference Format:
Qiong Wu and Chunyan Miao, 2013. Curiosity: From Psychology to Computation. *ACM Comput. Surv.* 0, 0, Article 0 (0), 26 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

What is curiosity? You are caught in a state of curiosity when you can't wait to read for the answer. Curiosity has been consistently recognized as a critical motivation that is associated with exploratory behaviors such as exploration, investigation, and learning [Kashdan et al. 2004]. It has been identified as a driving force for child development, scientific research, and educational achievement [Loewenstein 1994]. According to Kashdan [2010], curiosity benefits human beings at two levels: the individual level and the social level. At the individual level, curiosity is associated with individual growth, as an "innate love of learning and of knowledge...without the lure of any profit" [Loewenstein 1994]. At the social level, curiosity is an ingredient for enhancing interpersonal relationships through infusing energy and passion into social interactions.

Computational forms of curiosity have been developed to endow machines with desirable functions. For example, a curious design agent can arrange art exhibits to elicit the curiosity of their viewers and provide an aesthetically pleasing experience [Saun-

Author's addresses: Qiong Wu and Chunyan Miao, School of Computer Engineering, Nanyang Technological University, Nanyang Avenue, Singapore, 639798. E-mail: wuqi0005@e.ntu.edu.sg, ASCYMiao@ntu.edu.sg.

DOI:http://dx.doi.org/10.1145/0000000.0000000

ders and Gero 2001]; a curious exploratory agent achieves higher learning efficiency in unknown environments [Macedo and Cardoso 2005]; and a curious learning companion is able to detect interesting learning topics and artfully direct students' attention to them [Wu et al. 2012]. In this work, we attempt to offer a review on computational curiosity, with the aim to analyze the current research development and envision future trends.

From a machine learning perspective, curiosity has been proposed as algorithmic principles to focus learning on novel and learnable regularities in contrast to irregular noises [Schmidhuber 1991a; Oudeyer and Kaplan 2004; Pang et al. 2009; Karaoguz et al. 2011]. These principles make machines "explorative" and allow machines to have "the desire to improve the model's knowledge about the world" [Schmidhuber 1991c; 2006]. These principles have shown success in speeding up learning and building unsupervised developmental robotics [Schmidhuber 2006; Oudeyer and Kaplan 2004].

Human beings, unfortunately, are often more complex than what a machine learning model can describe. Most machine learning algorithms are strictly utilitarian. The machine always remains enthusiastic to learn and mostly assumes that there is something to be learned. On the other hand, a person may behave less rationally. For example, a person may lack the will to learn even when ample opportunities for learning exist. Alternatively, a person may be curious about a stimulus but realize that little can be learnt from it due to a lack of knowledge or adverse environments (e.g. information blocked by a firewall). In order to provide a more complete model of human cognition and design artificial companions that can elicit curiosity (e.g. for pedagogical purposes), we believe it is beneficial to go back to the research in psychology to understand how human curiosity can be aroused.

Research in psychology suggests that curiosity can be externally stimulated by various factors, such as novelty, conflict, uncertainty, and complexity, etc. [Berlyne 1960]. Each of these factors characterizes a different condition where curiosity can potentially be elicited. For example, novelty is induced by something new, whereas surprisingness occurs when an expectation based on previous knowledge is violated by the actual outcome. Several curiosity stimulating factors have been considered in computational curiosity. Saunders and Gero [2001] modeled novelty in their curious design agent; Macedo and Cardoso [2005] incorporated novelty, surprise, and uncertainty into their explorative agent; and Wu et al. [2012] considered novelty, uncertainty, conflict, and complexity in their curious pedagogical agent. The explicit consideration of different curiosity stimulating factors makes the curious agents more human-like, more efficient in exploration [Macedo and Cardoso 2005], and more reactive in collaborative learning environments [Wu et al. 2012]. Hence, we believe that the explicit consideration of various curiosity stimulating factors can enrich the current understanding of computational curiosity and open up new directions for future development.

In this paper, we attempt to answer the following research questions:

(1) What is the current common understanding on human curiosity in psychology?
(2) What is the current research progress on computational curiosity?
(3) What are the potential research directions for computational curiosity?

The paper is organized as follows. Section 2 provides a brief review on psychological studies of human curiosity and offers some fundamental guidelines for modeling computational curiosity in humanoid agents. Along this review of psychology, we propose a two-step process of curiosity arousal that serves as a psychological viewpoint to study its computational implementations. Subsequently in Section 3, we review existing computational models of curiosity in view of the two-step model and create taxonomy for them. Four key research issues regarding the theoretical models of computational curiosity are summarized in Section 4. In Section 5, the review is concluded by pointing out ten important research areas where computational curiosity can be infused to achieve significant innovations.

2. A BRIEF REVIEW ON HUMAN CURIOSITY

Research on human curiosity has attracted scientists' interest since back in the 1960s, when researchers attempted to uncover the underpinnings of curiosity [Loewenstein 1994]. The effort continues to date with researchers focusing on interpreting relationships between curiosity and human well-being [Kashdan 2010]. In this section, we briefly review psychological studies on human curiosity while answering two questions: (1) what are the underpinnings of curiosity, and (2) what is its arousal mechanism. Knowledge from the psychological research on curiosity can provide insights into modeling curiosity computationally.

2.1. Underpinnings of Curiosity

In human beings, curiosity is closely related to cognition, emotion and behaviors. It underlines human cognitive development, aesthetic appreciation, and interpersonal relationships. In this section, we review the literature in psychology on the underpinning of curiosity and the benefits of curiosity for human beings.

2.1.1. Categories of Curiosity. Most psychologists believe that curiosity is an intrinsic motivation driving the cognitive development of both humans and animals alike. Berlyne [1960] categorized curiosity along two spectrums: (1) from perceptual curiosity to epistemic curiosity, and (2) from specific curiosity to diversive curiosity. Perceptual curiosity, which resides in the lower level of cognition, stems from the senses of both animals and humans (e.g. senses of touch, vision, taste, etc.). It is defined as "a drive that is aroused by novel stimuli and reduced by continued exposure to these stimuli" [Loewenstein 1994]. Epistemic curiosity, referred to as "an appetite for knowledge", is related to the higher level of cognition and believed to be a distinctive human feature. While perceptual curiosity and epistemic curiosity are defined along the lines of "lower" and "higher" levels of cognition, specific curiosity and diversive curiosity are distinguished by the possibility of curiosity having a "direction". Specific curiosity is aroused by a particular piece of information. Diversive curiosity is a general drive to seek information with no specific direction and is predominantly employed to relieve boredom.

Berlyne's cognitive account for curiosity has been a theoretical foundation for the majority of recent studies. Litman and Spielberger [2003] agreed with Berlyne that there is a salient difference between diversive curiosity and specific curiosity, and conducted experimental analysis to provide scales for measuring both concepts. They further concluded that diversive curiosity and specific curiosity, as well as perceptual curiosity and epistemic curiosity, are "substantially correlated". Spielberger and Starr [1994] associated diversive curiosity with "low-level" stimuli and specific curiosity with "high-level" stimuli. However, Schmitt and Lahroodi [2008] disagrees with the notion that curiosity can be diversive. They argued that curiosity can only be specific towards its possessor and objects that cause the curiosity. Instead of diversive curiosity, they referred to the generic desire for knowledge as "inquisitiveness".

Nevertheless, a general consensus among the psychologists points to the close relationship between curiosity and cognition, as a drive to explore novel stimuli or an appetite for knowledge.

2.1.2. Curiosity-Related Emotions. Curiosity is related to emotional constructs. In an early account of curiosity by James [1950], curiosity is viewed as an instinctual or emotional response closely related to fear. He believed that curiosity motivates organisms to actively explore their environment, whereas fear tends to turn the organisms away from the risks induced by unfamiliarity. Berlyne's branch of research, referred to as "drive theory" by Loewenstein [1994], is based on the assumption that curiosity produces an unpleasant sensation that is reduced by exploratory behavior. Loewenstein believed that rather than serving a purposive end, the primary objective of satisfying one's curiosity is to induce pleasure.

Wundt [1874] introduced the theory of "optimal level of stimulation", which serves as a general rule postulating the relationships between stimulus intensity and hedonic tone. Based on this, Berlyne [1960] proposed that there is a need of "intermediate arousal potential" for curiosity to be aroused. This theory demonstrates that too little stimulation can result in boredom or inaction, while too much stimulation may result in aversion or withdrawal. Only when the level of stimulation levels is optimal and pleasurable can exploratory behavior occur.

From the above discussion , it can be seen that curiosity is closely related to emotional constructs such as "fear", "pleasure", "boredom" and "anxiety". The decision on whether to explore or to avoid a stimulus is driven by emotional comfort and results in behaviors that regulate emotional states.

2.1.3. Curiosity-Related Behaviors. The most salient expression of curiosity is through exploratory behaviors, by which curiosity can be satisfied. Berlyne [1960] defined two levels of exploratory behaviors, one associated with the perceptual level of curiosity, and the other associated with the epistemic level of curiosity. At each level, the exploratory behaviors can take many forms.

At the perceptual level of curiosity, Berlyne divided exploratory behaviors into three categories according to the nature of responses. He referred to the exploratory behaviors as *orienting responses* if they consist of changes in posture, orientations of sensory organs, or states of sensory organs. The second category of exploratory behaviors is associated with *locomotion*, such as approaching or withdrawing from the stimuli. When an exploratory behavior causes changes in external objects, through manipulation or otherwise, it is called an *investigatory response*.

At the epistemic level of curiosity, Berlyne also defined three categories of exploratory behaviors. The first category is *observation*, which places the subject in contact with external situations that can nourish the pertinent learning process. The second category is *thinking*, which refers to "productive" and "creative" thinking, rather than "reproductive thinking" that only calls up memories to determine how problems should be handled. The last category is *consultation*, which exposes an individual to verbal stimuli issued from other individuals, including questioning and reading, etc.

Through exploratory behaviors, cognitive growth is achieved by "creative thinking" and emotional states are regulated towards a feeling of "pleasure".

2.1.4. Benefits of Curiosity. Current research indicates that curiosity can contribute to human well-being at two distinct levels: individual level and social level.

At the individual level, curiosity provides an emotional motivation for self development. It is the driving force for child development as well as an important spur for educational achievement [Loewenstein 1994]. Also, literature in psychology indicates a close relation between curiosity and aesthetics, humor, and fun. According to Berlyne [1960], these behaviors have common motivational factors and are reinforced by common sources of reward.

At the social level, curiosity can enhance interpersonal relationships. By studying the role of curiosity in conversations, Kashdan et al. [2011] suggested that curiosity can build social bonds by promoting behaviors such as engagement, responsiveness, and flexibility. These are desirable behaviors for developing interpersonal relationships and building intimacy. Their findings indicate that curiosity is uniquely related to the development of interpersonal closeness with strangers [Kashdan et al. 2004].

2.2. Arousal of Curiosity

In this section, we review the literature in psychology on the arousal mechanism of curiosity and provide insights into possible ways of implementing curiosity in artificial beings.

According to Loewenstein [1994], existing psychological theories on human curiosity can be divided into three categories: *incongruity theory*, *competence theory*, and *drive theory*. Incongruity theory holds on to the idea that curiosity is evoked by violation of expectations [Hebb 1949]. Competence theory views curiosity as an intrinsic motivation to master one's environment [White 1959]. Drive theory believes in the existence of a curious drive, either primary (homeostatic generated in a similar way as hunger) or secondary (externally generated by stimuli). As commented by Loewenstein, the incongruity theory and the competence theory suffer from the same deficiency that both fail to offer a comprehensive account of curiosity. Hence, in this work, we focus on the drive theory and adopt Berlyne's interpretation of curiosity.

According to Berlyne [1960], traditional psychological research concentrated on problems of *response selection*, which studies what response humans will make to one standard stimulus at a time. However, curiosity deals with a different problem from response selection, which is referred to as *stimulus selection*. Stimulus selection discusses when several conspicuous stimuli are introduced at once, to which stimulus humans will respond. Berlyne has conducted extensive experimental studies to understand the process of stimulus selection and discovered a set of collative variables that govern this process.

2.2.1. Collative Variables. Collative variables, according to Berlyne [1960], refer to the external factors that govern various forms of stimulus selection. There are four major collative variables, viz., *novelty*, *uncertainty*, *conflict* and *complexity*. Berlyne named these factors collative variables because the evaluation of each variable involves an analysis of similarities and differences between elements in a stimulus pattern. In the following part of this section, the major collative variables are reviewed.

Novelty denotes something new. In accordance with the time when a stimulus has been experienced by an organism, novelty can be divided into *short-term*, *long-term*, and *complete* novelty. A stimulus can either be completely new to an organism, or it could have been experienced within just the last few minutes. In the former case, it is called complete novelty, and in the latter case, it is called short-term novelty. The intermediate case where a stimulus has not been experienced for a period of time (usually days) is called long-term novelty. Based on whether a stimulus possesses qualities that has been perceived before, it is classified into *absolute novelty* or *relative novelty*. If the stimulus does not have any previously perceived quality, then it is absolute novelty; otherwise, it is relative novelty.

Based on these observations, Berlyne introduced three criteria to measure novelty, i.e., novelty is inversely related to (1) how often the stimuli have been experienced before, (2) how recently the stimuli have been experienced, and (3) how similar the stimuli are to previously experienced ones.

Novelty is often accompanied by other properties, each of which may have different influence on exploratory behaviors. Berlyne listed them as supplementary variables to novelty, which include change, surprisingness, and incongruity. A **change** of the stimulus in question may induce some priority in exploratory directions. **Surprisingness** arises when there is a stimulus that induces an expectation and a later stimulus that contradicts the expectation. **Incongruity** is somewhat different from surprisingness (e.g., the statement that the Earth is round is a surprise to people who are accustomed to the concept that the Earth is flat), and indicates an expectation that is not met by the same stimulus (e.g., a person who is used to receiving gifts on his birthday before may find the experience incongruent when nobody remembered to send him gifts on his birthday this year).

Uncertainty arises when an organism has difficulty selecting a response to a stimulus. Berlyne adopted information theory to quantify uncertainty. He proposed to measure the uncertainty caused by a stimulus with the following steps: (1) draw up a list of stimuli that might occur (as a response to the stimulus in question), (2) partition them into classes, and (3) assign a probability to each class. The probability p_i of each class denotes the competing strength of each possible response, and Shannon's entropy $H = -\sum_{i=1}^{n}(p_i \times log_2(p_i))$ denotes the degree of uncertainty.

Conflict occurs when a stimulus arouses two or more incompatible responses in an organism. A response can be incompatible with one another in several ways. Firstly, some responses may be innately antagonist to each other. For example, no organism can move forward and backward at the same time. Other responses may initially be capable of performing together but become incompatible through learning. For example, we seldom frown when shaking hands. The third reason for incompatible responses may be attributed to the limitation of an organism's ability to multi-task. For example, it would be considered as an outstanding ability for a person to read two books at the same time.

Berlyne proposed four criteria for measuring conflict. Conflict is positively related to (1) the nearness to equality in the strengths of competing responses, (2) the absolute strengths of competing responses, (3) the number of competing responses, and (4) the degree of incompatibility between competing responses.

Complexity roughly refers to the variety or diversity in a stimulus pattern. Three most obvious properties that determine the degree of complexity are: (1) the number of distinguishable elements in a stimulus, (2) the dissimilarity between these elements, and (3) the degree to which several elements are perceived and responded to as a unit.

All the collative variables are interrelated. Berlyne postulated that novel stimuli can lead to conflict when a new pattern of stimulus can be sufficiently similar to several familiar stimulus patterns, which may induce many incompatible responses. Also, novelty has a connection with complexity. For example, a stimulus with short-term novelty will have a higher degree of temporal complexity than purely repetitive patterns. Alternatively, a higher degree of synchronous complexity may induce a higher degree of relative novelty. Moreover, complexity is also associated with uncertainty and conflict. For example, complex patterns with more parts can assume a larger amount of alternative forms and the degree of uncertainty can be affected by the number of classes in these alternative forms; each subelement of a complex figure may induce different responses that are incompatible with each other, leading to a higher level of conflict.

In summary, collative variables are properties in a stimulus that describe curiosity stimulating conditions. According to Berlyne, they are all eminently quantitative properties that exist plainly in varying degrees. The quantitative property of collative variables makes them potential candidates for measuring the stimulation level of curiosity. Most of the computational models of curiosity have considered at least one of the collative variables for the determination of stimulation level (a review of the same will be provided in Section 3).

2.2.2. Intermediate Arousal Potential. The existence of a stimulus does not necessarily result in curiosity. The arousal of curiosity depends on the appropriate level of stimulation that can be induced by a stimulus.

In the 1870s, Wundt [1874] introduced the concept of "optimal level of stimulation" and postulated an inverted "U-shape" relationship between the stimulation level and the hedonic value, referred to as the **Wundt curve** (Figure 1). It is a general rule stating that many forms of stimulation are pleasant at medium intensities and become

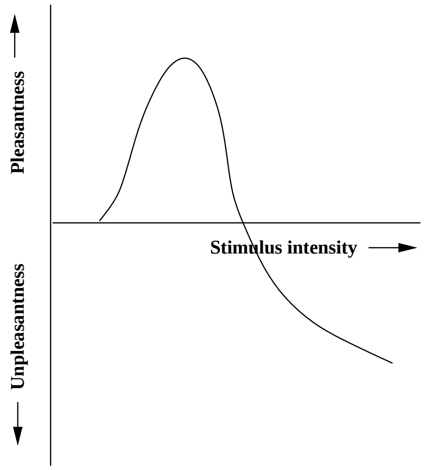

Fig. 1. The Wundt curve [Berlyne 1960].

unpleasant when their intensities are too high. Based on Wundt's theory and other experimental results from the literature, Berlyne formed the theory of "intermediate arousal potential", where too little stimulation results in boredom, too much stimulation results in anxiety and only intermediate stimulation results in curiosity. The two ends of the spectrum in the Wundt curve reflect two rules in stimulus selection: *Avoidance of Boredom* (**AoB**) and *Avoidance of Anxiety* (**AoA**).

To summarize, from the psychological point of view, the arousal process of curiosity can be abstracted into a **two-step model**, which offers a uniform angle for examining and analyzing existing computational models of curiosity.

— **Step 1**: evaluation of the stimulation level based on collative variables.
— **Step 2**: evaluation of the curiosity level based on the principle of intermediate arousal potential.

3. COMPUTATIONAL MODELS OF CURIOSITY

In this section, we review and analyze the existing computational models of curiosity from the perspective of the proposed two-step model. A general appraisal process of computational curiosity is summarized in Figure 2. The input stimuli are data samples that are perceived by agents and can trigger the appraisal process of computational curiosity. In step 1, the appraisal process first evaluates the level of stimulation in the stimuli based on the collative variables considered in each model. Some models adopt a single collative variable to determine the stimulation value (e.g., [Saunders and Gero 2001]), while others aggregate multiple collative variables to derive the stimulation value (e.g., [Macedo and Cardoso 2005; Wu et al. 2012]).

In step 2, the level of curiosity is evaluated through a mapping from the stimulation value to the curiosity value. Some models follow the principle of "intermediate arousal potential" by explicitly simulating the Wundt curve, which represents a nonlinear mapping from stimulation to curiosity (e.g., [Saunders and Gero 2001; Merrick et al. 2008]). These models accommodate both AoB and AoA in their stimulus selection approaches. Other models simply use the stimulation value as the curiosity value (e.g., [Schmidhuber 1991b; Oudeyer and Kaplan 2004; Macedo and Cardoso 2001]) or assume a positive correlation between the stimulation level and the curiosity level (e.g., [Wu et al. 2012]). Some of the models which directly use the stimulation value as the curiosity value consider the principle of "intermediate arousal potential" when

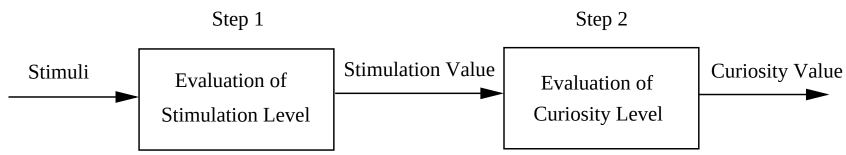

Fig. 2. The General Appraisal Process of Computational Curiosity.

determining the stimulation value (e.g., [Schmidhuber 1991b; Oudeyer and Kaplan 2004]), which also accommodate both AoB and AoA. The rest of the models which assume positive correlations between stimulation and curiosity only support AoB but not AoA because high stimulation can lead agents into anxiety (e.g., [Macedo and Cardoso 2001; Wu et al. 2012]).

A taxonomy for the existing models of computational curiosity is provided in Table I. The models are classified into six categories according to the collative variables used for the evaluation of stimulation in Step 1. These six categories include novelty, surprise, uncertainty, conflict, change, and complexity. Within each category, the models are again classifies into two sub-categories, according to whether the principle of "intermediate arousal potential" is followed in the models. The first sub-category of models follows this principle and supports both AoB and AoA in stimulus selection, while the second sub-category only supports AoB.

3.1. Models based on Novelty

As mentioned before, Berlyne associated three criteria with the evaluation of novelty: the degree of novelty is inversely correlated to (1)how often the stimuli have been experienced before (*frequency*), (2) how recently the stimuli have been experienced (*recency*), and (3) how similar the stimuli are to previously experienced ones (*similarity*). From this description, it can be observed that the appraisal of novelty is time-dependent and involves a comparison between the current stimuli and previous experiences.

3.1.1. Similarity-based. To explore the possibility of machines capable of generating creativity, Saunders and Gero [2001] developed a computational model of curiosity for intelligent design agents, focusing on the appraisal of novelty. The goal of such an agent is to evaluate the interestingness of creative design patterns based on individual experiences, where curiosity is the key for the evaluation of interestingness and the selection of design patterns [Saunders and Gero 2001]. In their model, each creative design pattern (generated by a design generator) can be considered as a stimulus and the stimuli experienced in the past are stored in a "conceptual design space". This "conceptual design space" is modeled by a Self-Organizing Map (SOM), representing the forms that the agent has experienced often enough to learn.

If we view their model from the perspective of the two-step model, in Step 1, the evaluation of stimulation is governed by the degree of novelty in a stimulus. The evaluation of novelty is achieved by comparing the stimulus (a creative design pattern) with past experiences (the SOM representations), and is defined by the complement of the typicality measure. This implementation is in line with Berlyne's third criteria for measuring novelty: novelty is inversely proportional to similarity (typicality).

In Step 2, Saunders and Gero adopted the full principle of intermediate arousal potential and explicitly modeled the Wundt curve as a non-linear function. This function is defined by the difference of two sigmoid functions as follows [Saunders 2002]:

$$H = R - P$$

Table I. Taxonomy of the Computational Models

Collative Variables	Avoidance of Boredom & Avoidance of Anxiety		Avoidance of Boredom	
Novelty	Similarity & Simulation of the Wundt curve	Saunders & Gero[01] Maher et al.[08]	Similarity	Macedo & Cardoso[99]
			Frequency& Similarity	Ogino et al.[06]
			Recency, frequency &Similarity	Wu et al.[12]
Surprise	Prediction error in zero sum games	Schmidhuber[99]	Prediction error	Schmidhuber[91a] Uğur et al.[07]
	Information gain	Storck et al.[95]	Improbability	Macedo & Cardoso[99]
Uncertainty	-		Entropy	Macedo & Cardoso[95]
			No. of elements that contains uncertainty	Wu et al.[12]
Conflict	-		No. of elements that contains conflict	Wu et al.[12]
Change	Spread adjustment in Gaussian	Karaoguz et al. [11]	-	
Complexity	Prediction improvement between successive situations	Schmidhuber [91b]	No. of elements that reflects complexity	Wu et al.[12]
	Prediction improvement between similar situations	Oudeyer & Kaplan[04]		
	Compression improvement	Schmidhuber [06]		
	Discriminability difference between two spaces	Pang et al.[09]		

$$R = \frac{R_{max}}{1 + e^{-\rho_R(novelty - R_{min})}} \quad (1)$$

$$P = \frac{P_{max}}{1 + e^{-\rho_P(novelty - P_{min})}}$$

where H is the curiosity value, R, P, R_{max}, P_{max}, ρ_R, ρ_P, R_{min} and P_{min} are the reward, punishment, maximum reward, maximum punishment, slope of the reward sigmoid functions, slope of the punishment sigmoid functions, minimum novelty to be rewarded and minimum novelty to be punished, respectively. Based on this curiosity value, the agent can report evaluation of the interestingness of a design pattern or how curious it is about the design pattern.

Later, Merrick and Maher applied this model of curiosity to controlling the behaviors of Non-Player Characters (NPCs) and reconfigurable robots [Merrick et al. 2008; Merrick and Maher 2009; Maher et al. 2008; Merrick and Huntington 2008; Merrick 2008].

For NPCs, the goal of infusing curiosity is to achieve creative behaviors, namely, behavioral diversity [Merrick et al. 2008; Maher et al. 2008]. These works are rooted in a motivated reinforcement learning framework, by which NPCs can learn the state-action policy in a virtual environment. Here, each event E is regarded as a stimulus and is defined as the change between the current state and the previous one. All previously experienced events are stored in a SOM structure. A habituation layer of neurons is connected to the clustering layer and a habituation function is used to compute novelty. The final curiosity value is obtained by feeding the novelty level into the simulated Wundt curve (Equation (1)). The curiosity value is used as an intrinsic reward to update the policy π that maps states to actions. This curiosity reward can eventually lead to creative behavior, because each new interesting event can result in a perturbation of the existing behavioral patterns and form reward signals for the learning process.

For reconfigurable robots [Merrick and Huntington 2008], the role of curiosity is to direct the robots' attention to reconfigurations in their structures. This is an important skill for robots to learn new behaviors in response to structural changes. Similar to the implementation of curiosity for NPCs, an event can be considered as a stimulus and defined as the change between the current sensation and the previous one. Following a similar algorithm used for curious NPCs, curiosity reward is generated for robots. With this curiosity reward, the robots are self-motivated to explore changes in their structures and develop new behaviors.

In the application domain of exploring unknown environments (e.g. planetary exploration, map-building of interiors, etc.), Macedo and Cardoso [1999] proposed a model of curiosity for intelligent agents to simulate human-like exploratory behavior. Along their research, they gradually introduced novelty, surprise, and uncertainty into their computational model of curiosity. Here, we first look at their model of novelty.

Macedo and Cardoso's model relies on graph-based mental representations of objects. Each object can be regarded as a stimulus. In Step 1, the level of novelty regarding a stimulus (object) is measured based on the error correcting code theory of Hamming. Three steps are considered: 1) representing each graph (describing an object) in a common shape matrix 2) extracting the numerical code from the matrix representation of each graph, and 3) computing the Hamming Distance. Here, novelty is defined as the minimum Hamming Distance from the stimulus to the ones that have been experienced before. Novelty determines the stimulation value. In Step 2, the stimulation value is directly used as the curiosity value. The system only supports AoB because a higher level of curiosity increases the chance of an object to be explored.

For both models by Saunders and Gero [2001] and Macedo and Cardoso [1999], regardless of the implementation, i.e., atypicality in SOM and hamming distance in graph-based mental representations, the model of novelty reflects a comparison between the current stimuli and previous experiences (in other words, similarity).

3.1.2. Time and Similarity-based. In some later models of curiosity, the factor of time, as another dimension in addition to similarity, is considered for measuring novelty.

Ogino et al. [2006] addressed lexical acquisition problems for robots using computational curiosity to associate visual features of observed objects with the labels (for objects) that are uttered by a caregiver. In their work, each object can be interpreted as a stimulus for the robot. Visual features of each object are represented by a SOM. In Step 1, stimulation value is determined by novelty. The novelty of each object is calculated based on two types of saliency: habituation saliency, reflecting the temporal saliency; and knowledge-driven saliency, reflecting the similarity saliency.

The habituation saliency is characterized by habituation and inversely related to the frequency of observation of a visual feature. The habituation saliency S_1^i is updated as follows:

$$S_1^i(t-1) + \Delta S_1^i(t-1) \tag{2}$$

$$\Delta S_1^i(t) = \frac{\alpha(1 - S_1^i(t)) - \beta S_1^i(t) I^i(t)}{\tau} \tag{3}$$

where α is a constant that characterizes spontaneous recovery, β is a constant that determines the rate of habituation, τ is the time constant, and I^i is the activation level of the i-th neuron in the visual feature map.

The knowledge-driven saliency is characterized by the acquired knowledge, where more saliency is given to visual features that are not associated with other labels. The knowledge-driven saliency S_2^i is given as follows:

$$S_2^i = 1 - sigmoid(\omega_{L \to i}) \tag{4}$$

$$sigmoid(w) = \frac{1}{1 + e^{-\alpha(\omega - \theta)}} \tag{5}$$

where $\omega_{L \to i}$ represents the index of the familiarity of the i-th visual feature neuron, α is the parameter that determines the rate of rise of sigmoid function and θ is a threshold.

The product of the two saliency values represents the overall level of stimulation, which is directed used as curiosity value (in Step 2) and is represented by:

$$C^i(t) = (S_1^i(t) + c_1) \times (S_2^i(t) + c_2) \tag{6}$$

where c_1 and c_2 are constants.

The robot chooses to learn about the object with maximum curiosity value, which only considers AoB in stimulus selection. The infusion of curiosity helps accelerate the robot's learning.

Wu et al. [2012] attempted to infuse human-like curiosity into a virtual companion to detect potentially interesting learning objects for users and help them avoid the feeling of being lost in virtual world-based learning environments. Their model is built based on Berlyne's theory and considers four collative variables that can stimulate curiosity: novelty, uncertainty, conflict, and complexity. In this section, we first look at their model of novelty, which combines time-based measures (frequency and recency) and similarity-based measures.

The curious companion maintains two sources of information: world knowledge, which is defined by experts and embedded in the virtual world, and user knowledge, which is updated by users in real time. Both types of knowledge are represented by Concept Maps (CMs). In the CM for world knowledge, each concept is embodied by a virtual object, and each relationship between concepts specifies a knowledge point that can be learnt by users through interactions with the virtual objects. The CM for user knowledge represents the human learner's current understanding.

Here, each virtual object embedded with learning contents can be considered as a stimulus. In Step 1, novelty of a stimulus is determined by three factors: frequency, recency, and similarity. Frequency (how many times a virtual object has been met in a time period) and recency (how recently a virtual object has been met) are time-related variables that can be retrieved from user's behavior history. Similarity is an integer determined by the number of knowledge points learnt by the learner correctly, through the comparison between world knowledge and user knowledge. Novelty is calculated based on a fuzzy combination of all the three factors using a Fuzzy Cognitive Map (FCM) based approach. In Step 2, a linear relationship is assumed between stimulation and curiosity, and the curiosity value is positively proportional to the level of novelty (specified by the weight of FCM). The combination of novelty and the other three collative variables (which will be introduced in subsequent sections) determines the final level of curiosity. The agent will select the virtual object with the highest curiosity value for the user to explore, which only supports AoB in stimulus selection.

3.2. Models based on Surprise

Two interpretations for surprise exist in the literature of computational curiosity. The first one interprets surprise as the difference between an expectation and the real outcome. Prediction error matches well with this interpretation and has been utilized in many curiosity models to measure the level of surprise [Schmidhuber 1991a; Barto et al. 2004; Schmidhuber 1999; Uğur et al. 2007]. The second interpretation describes surprise as the degree of not expecting something. Storck et al. [1995] modeled this type of surprise using the information gain before and after an observation, while Macedo and Cardoso [2001] proposed another measure using improbability.

3.2.1. Prediction Error-based. Schmidhuber [1991a] introduced artificial curiosity into model building control systems. The goal of such a control system is to learn the input-output mapping in a noisy environment. Curiosity is infused to give the control system an intrinsic desire to improve the model's knowledge about the world. This is realized by introducing an additional reinforcement unit on top of the controller, which rewards actions that cause high prediction errors. The prediction error inherently measures the degree of mismatch between belief and reality, which is in line with the first interpretation of surprise. In Step 1, the stimulation value is determined by surprise (prediction error) and it is directly used as the curiosity value (in Step 2) to form decisions (on rewarding the system). This mechanism (rewarding the system with high surprise/curiosity) encourages certain past actions to be carried out again in order to repeat situations similar to the mismatched ones. Hence, the system will always direct its attention toward something that is unknown and therefore avoid boredom. This working mechanism only supports AoB in stimulus selection. Barto's theory of intrinsically motivated reinforcement learning is implemented based on a similar principle [Barto et al. 2004; Singh et al. 2004].

In a later variation, Schmidhuber [1999; 2002] worked on exploring the space of general algorithms that can automatically create predictable internal abstractions of complex spatio-temporal events. In this model, both AoB and AoA are accommodated. Curiosity is interpreted as the ability of the system to focus on interesting things by losing interest in the overly predictable (boring) or the overly unpredictable (anxiety inducing) aspects of the world. To achieve this goal, the system is realized by two intrinsically motivated agents playing zero-sum games. The two agents can bet in advance on whether a prediction is true or false. If they bet on different outcomes, the system will check who is right. The winner gets rewarded by receiving the other's bids whereas the loser loses its bids due to surprise (error in prediction). Hence, both agents are motivated to lure the opponent into agreeing on computation sequences that will surprise the other one. However, a surprised module will eventually adapt and in turn, cause the other agent to lose a source of reward. In this way, the system as a whole is motivated to shift the exploration focus and reveal the unknown yet predictable regularities.

Another example of prediction error-based implementation of surprise is given by Uğur et al. [2007]. In situations where a robot physically interacts with the environment to explore and learn, assuming the traditional reinforcement learning method is applied, the robot may require a large number of interactions to learn even simple tasks. This could eventually damage the robot. To address this problem, Uğur adopted a Support Vector Machine (SVM) to learn the perception and action mapping in the robot, where curiosity is introduced to select interesting training data for SVM to reduce the total number of training data required. In Step 1, the stimulation level is determined by surprise, which is measured by a sample's distance to the hyperplane (which separates two classes) in the feature space. In Step 2, the stimulation value is directly used as curiosity value for decision making. The system supports AoB because it can be driven away from boredom: only if the distance is smaller than a fixed threshold (curiosity value is high), the sample is considered interesting and sent for learning.

3.2.2. Probability-based. An extension of Schmidhuber's curiosity principle to non-deterministic environments was done by Storck et al. [1995]. The goal of their system is to learn the model of a non-deterministic Markov environment where each state-action pair $(S(t), a(t))$ may result in different next states $(S(t+1))$ probabilistically. The goal of introducing curiosity into this learning system is to actively search for interesting training examples that can maximize expected learning improvement. In step 1, the stimulation value is determined by surprise, which reflects "the degree of not expecting something". They adopt a probabilistic way of measuring surprise, which is the Kullback-Leibler distance between the agent's belief distribution before and after making an observation. This value of information gain reflects the learnability of the exploration space. Exploration areas where little information gain can be achieved are either too predictable (well-learnt) or too unpredictable (inherently unlearnable). Hence, the information gain (surprise) accommodates both AoB and AoA in stimulus selection. In step 2, the surprise value is directly used as the curiosity value, which acts as rewards for the reinforcement learning of state-action pairs. In this way, curiosity drives the system away from the exploration space where little information gain can be achieved, i.e., areas that are either too boring or anxiety-inducing.

Macedo and Cardoso [1999; 2001; 2004; 2005], mentioned in Section 3.1.1, also modeled human forms of surprise in artificial agents. Surprise serves as an intrinsic drive to direct the agent's exploration in unknown environment. Their model relies on graph-based mental representations of objects, where each object is considered as a stimulus. The surprise level is defined by the degree of not expecting a stimulus (object), and is implemented as the improbability of existence of the stimulus. The improbability is evaluated based on probability theory: the degree of expecting an event X is given by its probability $P(X)$, and accordingly, the improbability of X , denoted by $1 - P(X)$, defines the degree of not expecting X (surprise). In the case of an object Obj_k comprising several components, the probability of expecting an object is calculated by the mean value of conditional probabilities of its n constituent parts, where each part is individually computed based on Bayes's formula, given by

$$P(Obj_k) = \frac{\sum_{l=1}^{n} P(Obj_k^l | Obj_k^1, Obj_k^2, \cdots, Obj_k^{l-1}, Obj_k^{l+1}, \cdots, Obj_k^n)}{n} \tag{7}$$

where $P(Obj_k)$ is the probability of expecting the k^{th} object.

However, surprise in this model is not treated as a stimulating factor for curiosity. Hence, the surprise value is not mapped to the curiosity value.

3.3. Models based on Change

In robot sensory-motor learning, Karaoguz et al. [2011] proposed a model of curiosity to direct the robot's attention to changes in the environment. The aim of the model is not to measure the amount (or level) of changes in a stimulus, but rather to focus on how to react when changes occur. Hence, the system does not provide evaluation for the level of stimulation or the level of curiosity (induced by change), but offers mechanisms to redirect its attention to changes.

For the robot, the focus of attention is determined by a Gaussian distribution that governs the sampling of training examples over the laser motor space. The center of the Gaussian distribution is determined by the mean value of the last N samples added to the mapping, and directs the attention of the robot to areas where new samples have been recently added. The spread of Gaussian distribution is related to the performance (success rate) of the mapping by

$$\sigma(t, s_r) = \sigma_0 + k_b t_b + k_f (t_f - s_r)^2 \tag{8}$$

where σ_0 is the baseline, t_b is the number of time steps since the last new sample was added, k_b is the boredom coefficient, t_f is the failure threshold, s_r is the success rate, and k_f is the failure coefficient.

It can be seen from Equation (8) that the attention spread is inversely correlated to the success rate. For example, the attention spread is wide when the success rate is low due to a high value of $k_f(t_f - s_r)^2$. This model can successfully redirect the robot's attention towards changes, because when a change happens, t_b will be reset to 0 and the high past success rate of that region will result in a narrow Gaussian distribution of samples. Newly-added links in unlearnable areas (noisy region) will fail on retest, which can drive the success rate for that region below the failure threshold t_f and increase the width of sampling distribution, redirecting the system away from that region. Hence, the system has a drive towards regions where existing mappings are not making accurate predictions (AoB) and a drive away from regions where no improvements can be obtained (AoA).

From the perspective of machine learning, one can argue that this model is a special case of earlier models based on predictors [Schmidhuber 1991a; 1991b], because a predictor always predicts that the next observation is similar to the current observation, which indicates no change. Once change is present, the predictor makes an error and updates its knowledge about the world.

3.4. Models based on Uncertainty

Uncertainty arises when there is no clear response to a stimulus [Berlyne 1960]. The entropy in information theory has been proposed to measure the degree of uncertainty in a stimulus. This measure has also been employed very often in computational implementations to realize uncertainty-based curiosity [Macedo and Cardoso 2005]. However, when stimulus is difficult to classify, entropy can hardly be used. In a virtual learning environment, Wu et al. [2012] proposed a rough measure of uncertainty in a stimulus (learning object) by counting the number of uncertain elements.

3.4.1. Information Theory-based. To improve an agent's exploration in unknown environments, Macedo and Cardoso [2005] introduced a measure of uncertainty, on top of novelty [Macedo and Cardoso 1999; 2001; 2004], into their model of curiosity. Based on the same system setup as presented in Section 3.1.1, Macedo and Cardoso [2005] argued that the desire to know or learn an object can be induced by both novelty and uncertainty, i.e., objects with at least some parts that are not yet known. Each object (stimulus) can contain known parts (without uncertainty) and uncertain parts. The known parts of an object are used to measure novelty by Hamming Distance(introduced in Section 3.1.1), whereas uncertainty is measured by the entropy of all uncertain parts, including analogical and propositional descriptions of the physical structure, and functions of the object. In step 1, the stimulation value is determined by the aggregation of novelty and uncertainty. In step 2, the stimulation value is directly used as curiosity value. The system adopts AoB in stimulus selection and chooses objects with highest curiosity value to explore.

3.4.2. Similarity-based. As mentioned before in Section 3.1.2, Wu et al. [2012] also considered uncertainty in their model of curiosity. For the virtual learning companion, uncertainty arises when it cannot verify a knowledge point (given by the user) with the expert knowledge. Hence, in Step 1, uncertainty is an integer calculated by using the world knowledge (expert knowledge) as reference and counting the number of knowledge points that are present in the user knowledge but are not present in the world knowledge. In Step 2, they assumed a positive correlation between stimulation and curiosity. The level of curiosity is set to be positively proportional to uncertainty (specified by the weight in the FCM-based approach).

3.5. Models based on Conflict

Conflict arises when a stimulus tends to induce two or more incompatible responses in an organism [Berlyne 1960]. For the virtual learning companion, Wu et al. [2012] defined conflict as the incompatibility between the human learner's understanding and the expert knowledge embedded in the virtual world. If a relationship from concept C_i to concept C_j in the user knowledge has a different label with the relationship from C_i to C_j in the world knowledge, then the user's understanding of this relationship is in conflict with the expert's. Hence, the measurement of conflict involves a comparison between the user knowledge and the world knowledge and is an integer calculated by counting the number of knowledge points having conflicting labels in both knowledge bases. In Step 1, the level of conflict in a virtual object (stimulus) is defined by the total number of conflicting relationships. In Step 2, the level of curiosity is positively correlated to conflict (specified by the weight in the FCM-based approach).

3.6. Models based on Complexity

In machine learning models of curiosity, complexity has been associated with the predictive power of predictive systems or compressive power of data compression systems [Li and Vitanyi 2008; Schmidhuber 1991b; 2006; Oudeyer and Kaplan 2004]. In virtual learning environment, complexity has been interpreted as how difficult a topic is to a student [Wu et al. 2012].

3.6.1. Prediction Improvement-based. In Section 3.2.1, we introduced Schmidhuber's early ideas on implementing curiosity by rewarding the system proportionally to surprise, i.e., prediction error [Schmidhuber 1991a]. However, this implementation only succeeded in guiding the system to avoid boredom (i.e., well learnt area) but not anxiety (i.e., inherently unpredictable area caused by noisy). Later, Schmidhuber refined the model to accommodate both AoB and AoA. He defined curiosity as a simple principle: to learn a mapping from actions to the expectation of future performance improvement [Schmidhuber 1991b; 1991c]. Instead of pure prediction error, reward is now generated according to the controller's prediction improvement, i.e., change in prediction error. In step 1, the complexity of a data to be learnt by the system (e.g., familiar data that is easy to learn or noises that are too difficult to learn) is determined by prediction improvement, which supports both AoB and AoA in stimulus selection. The prediction improvement is obtained by a confidence module, which evaluates the reliability of a prediction and can be realized by probability-based or error-based methods. In step 2, the stimulation value is directly used as curiosity value, which is adopted as intrinsic rewards for learning. In this way, the controller will choose actions (based on the delayed reward) to deliberately select training examples with easily learnable regularity. Hence, the infusion of curiosity can direct the system's attention away from the exploration space that is either too predictable or too unpredictable.

Schmidhuber [2006; 2009a; 2009b] formed his formal theory of creativity, by generalizing the simple principle of curiosity from predictors to data compressors. A 'beautiful' sensory data is one that is simple yet has not been fully assimilated by the adaptive observer, which is still learning to compress data better. The agent's goal is to create action sequences that can extend the observation history to yield previously unpredictable but quickly learnable algorithmic regularities. In other words, it is looking for data with high compressibility (reflected by curiosity value).

Schmidhuber's implementation of prediction improvement in nature is a comparison of prediction error between situations that are successive in time. This principle allows robots to avoid long periods of time in front of a television with white noise (completely unlearnable situations) because the error in prediction will remain large and the robots will be bored due to little prediction improvement. However, the principle is not robust enough in the alternation of completely predictable and unpredictable situations, because robots can get stuck here due to large prediction improvement. To cope with such problems, Oudeyer and Kaplan [2004; 2005] refined Schmidhuber's simple principle. Instead of comparing the prediction error between situations that are successive in time, compare the prediction error between situations that are similar. They proposed a model of curiosity that allows a robot to group similar situations into regions where comparison between situations is meaningful. The learning space is divided into regions and each region has an expert to make local predictions. Each expert computes the prediction improvement (curiosity value) locally and rewards its state-action pairs according to the prediction improvement. This works well in practice to handle problems such as robots stuck in situations with completely predictable and unpredictable sample data in alternation.

Another variation of using prediction improvement as curiosity drive has been proposed by Pang et al. [Pang et al. 2009]. This model is rooted in incremental Linear Discriminant Analysis (LDA). Here, curiosity is measured as the discriminability difference (residue) between the LDA transformed space and the original space. The infusion of curiosity can help the system actively search for informative examples to learn and improve the performance using fewer instances.

3.6.2. Similarity-based. In the model built by Wu et al. [2012], complexity is interpreted as the difficulty of a virtual object (stimulus) to the learner. The more knowledge points contained in a virtual object that are not new to a learner, the more complex (or difficult) it is to the learner. Hence, the complexity for a virtual object is an integer measured by the number of knowledge points (related to the virtual object) in world knowledge that have not been learnt by the learner. In step 2, Wu et al. assumed a positive correlation between stimulation and curisoity. Complexity, novelty (Section 3.1.2), uncertainty (Section 3.4.2), and conflict (Section 3.5) are aggregated by a FCM-based approach to calculate the level of curiosity. Once the learner is detected to be in a state of being lost, the agent will calculate a curiosity value for each virtual object in the learner's vicinity in the virtual world and prompt the learner to explore the virtual object with the highest curiosity value (potentially most interesting), which only considers AoB in stimulus selection.

3.7. Discussion

An interesting point worth noting is that, from the machine learning perspective, there are certain recurring principles underlying curiosity-inspired algorithms. The first group of recurring principles includes generating intrinsic curious rewards based on errors [Schmidhuber 1991a; Uğur et al. 2007] or Shannon's information [Scott and Markovitch 1989]. This group of principles can redirect learning to focus on the unknown. However, they fail to distinguish noise from the novel and learnable regularities. The second group of recurring principles includes generating intrinsic curious rewards based on error reduction [Schmidhuber 1991b; Oudeyer and Kaplan 2004], information gain [Storck et al. 1995], or compression improvement [Schmidhuber 2006]. This group of principles effectively addresses the above mentioned problem. They are able to guide learning to focus on easily learnable regularities and at the same time filter out noise. The second group of principles helps form the basis of Schmidhuber's theory of artificial curiosity, which shows success in speeding up learning and building unsupervised developmental systems [Schmidhuber 1991b]. Also, Schmidhuber believes that these principles make machines "creative" and intrinsically motivate machines to create action sequences that make data interesting, which forms the basis of his theory of artificial creativity [Schmidhuber 2006].

4. SUMMARY OF RESEARCH ISSUES IN COMPUTATIONAL CURIOSITY

Four key research issues are identified which deserve to be explored further to enhance the theoretical models of computational curiosity. These research issues are mainly concerned with designing agents that can demonstrate curiosity and perceive the level of curiosity exhibited by human beings.

4.1. Evaluation of Stimulation

Most of the current computational models of curiosity adopt a heuristic approach for the evaluation of stimulations. This is often done based on algorithmic characteristics and goals in machine learning. These models advance the performance of machine learning systems in learning and exploration. However, they have difficulties supporting humanoid agents to evaluate curiosity stimuli in complex environments such as computer-based education, e-commerce, and teleconference. Some models have started from a psychological theory and evaluated stimulation levels based one collative variable or a subset of collative variables. As of yet, how the collative variables affect the level of stimulation, individually or collectively, has not been studied. Both qualitative and quantitative analysis of collative variables can help form a deeper understanding of the working mechanism of computational curiosity, and provide a clear picture of how collative variables are related to performance changes in intelligent agents.

4.2. Evaluation of Curiosity

Existing computational models mostly assume a positive correlation between stimulation and curiosity, which may not be always true in human beings. Studies have been done on simulating the Wundt curve to map the level of stimulation to the level of curiosity, but the understanding of how this mapping affects the performance of human-like agents is still unclear. Moreover, in psychology, the curiosity zone is right next to the boredom zone and the anxiety zone. Hence, deeper studies are necessary to provide more proper mapping methods to allow human-like agents to avoid entering the boredom zone or the anxiety zone during exploration.

4.3. Curiosity-based Decision Making

One important issue with computational curiosity is the risk management in curiosity-based decision making. As curiosity often leads to explorations, sometimes the agent or human being may be exposed to the possibility of being harmed or causing undesirable consequences to others. Hence, curious machines should operate under the protection of proper risk management systems so that they will not harm themselves or others. Ethical boundaries should also be defined for curious machines so that they will not intrude on the privacy of the users or other agents.

4.4. Recognition and Response to Users' Curiosity

Currently, computational curiosity mainly focused on designing models for machines to demonstrate curiosity, which has been shown to be useful in improving their exploration and learning. Another key research direction for computational curiosity which has not been intensively studied is to understand the curiosity exhibited by the users. This topic is especially important for designing next-generation human-agent collectives in areas such as embodied conversational agents, computer-based education, and artificial companions. The ability of intelligent agents to recognize and respond to users' curiosity provides opportunities to enhance user experiences in human-computer interactions by avoiding boring and anxiety-inducing topics, or by stimulating users' curiosity appropriately. This research may also enrich the current understanding of human curiosity in psychology.

5. TEN RESEARCH AREAS THAT COMPUTATIONAL CURIOSITY WILL IMPACT

In this section, ten important research areas where computational curiosity could play significant roles are identified. We believe that computational curiosity is a novel dimension which can be infused into many important research areas and become a fertile ground for significant innovations.

5.1. Machine Learning

The close-knit relationship between curiosity and human learning has inspired many researchers to devise computational forms of curiosity for machine learning systems, with the expectation to enhance learning capability and potentially drive them to evolve into autonomous intelligent machines. The study of computational curiosity in machine learning systems has some overlap with other concepts such as "active learning" and "intrinsically motivated learning".

Active learning, as the name suggests, deals with making machines "active" by allowing the learning system to select actions or make queries that influence what data to be added into its training set [Cohn et al. 1996]. Active learning is especially useful when data are expensive or difficult to obtain. Curiosity can play a critical role in active learning by helping the system to determine which data are interesting. For example, Scott and Markovitch [1989] introduced a curiosity drive into supervised learning systems to actively select the most informative samples for learning. Here, the system is continually directed towards regions with highest uncertainty, which is also a general principle followed by many other active learning algorithms [Fedorov 1972]. Uğur et al. [2007] infused curiosity into a SVM-based learning system to select interesting training data, which significantly reduces the number of training samples required to learn. Similarly, Pang et al. [2009] introduced curiosity into a LDA-based learning system.

Intrinsically motivated learning advocates the development of “intrinsic motivations” for learning systems to achieve task-independent learning [Barto et al. 2004; Singh et al. 2004] or autonomous development [Oudeyer and Kaplan 2004]. These learning approaches are gaining increasing popularity among AI researchers [Baldassarre 2011]. Intrinsically motivated learning often takes root in a reinforcement learning framework, where intrinsic motivations act as intrinsically generated rewards that are to be maximized. In human psychology, curiosity is known to be one of the most important intrinsic motivations related to learning. Hence, curiosity has often been adopted in intrinsically motivated learning algorithms. For example, Barto et al. [2004] used Berlyne’s theory as the psychological foundations to develop his intrinsically motivated learning algorithm. Schmidhuber [1991b; 1999; 2009b] introduced artificial curiosity as intrinsic rewards for general learning algorithms. Oudeyer and Kaplan [2004] proposed an intelligent adaptive curiosity mechanism for intrinsically motivated robots.

5.2. Robotics

The role of curiosity in attention focus, self-development and autonomous learning has attracted intense attention in the field of robotics [Oudeyer and Kaplan 2004; Uğur et al. 2007; Merrick et al. 2008; Karaoguz et al. 2011].

With the attempt to design robots that can autonomously self-develop in a progressive manner, Oudeyer and Kaplan [2004] devised for them a mechanism that resembles human curiosity. This mechanism acts as an intrinsic motivation to motivate robots to explore into regions with new knowledge [Oudeyer and Kaplan 2007], and endows them with the ability to adapt to new environments without prior knowledge or manual adjustment. With a similar goal of designing robots that can self-develop without an explicit teacher, Ngo et al. [2012] applied Schmidhuber’s principle of curiosity into a robot arm that enables robots to learn skills through playing. Pape et al. [2012] applied the same into a biomimetic robot finger for learning tactile skills.

Traversibility affordance refers to the ability of robots to navigate through an environment with obstacles. This ability is highly dependent on the robot’s current location, orientation, and the shape of objects in the environment. In situations where robots physically interact with the environment to explore and learn, assuming traditional reinforcement learning methods are applied, even simple tasks such as avoiding objects may require a large number of trials. This increases the risk of the robot being damaged during the exploration. To address this problem, Uğur et al. [2007] simulated curiosity in robots to select informative training samples, which can significantly reduce the number of interactions required with minimal degradations in the learning process.

Another problem in robotics is self-directed reconfiguration. Reconfigurable robots can rearrange their modules to achieve different structures, behaviors, and functions. Instead of looking into how robots can adapt in an unstructured environment, reconfigurable robots focus on adaptation to changes in their own structures and changes of goals when the actuator or effector of the robot changes. Merrick and Huntington [2008] introduced curiosity into reconfigurable robots to select informative samples to learn, so that with fewer interactions, robots can still achieve better learning outcomes.

One of the most important sensory-motor problems for robots when interacting with an environment is to learn a mapping from the gaze space (the location of an object) to the reach space (the movement of arms to grasp the object) in response to changes in the environment (camera replacement or changes in the physical environment) without manual recalibration of the hardware. To address this problem, Karaoguz et al. [2011] devised a mechanism of curiosity that drives the exploration into learning spaces where a proper level of complexity is associated with a particular level of capability. With this mechanism, robots can concentrate on highly interesting areas that are neither fully explored nor pure noise. In addition, the mechanism can successfully direct robots' attention to regions where changes have occurred.

Exploration in extreme environments can be a dangerous task for humans. Artificial agents, especially in the form of robots, have been a good substitute for humans to undertake such tasks. Exploration of unknown environments has been an active research field in domains such as planetary exploration, meteorite searches in Antarctic, volcano exploration, and map-building of interiors, etc. [Moorehead et al. 2001; Burgard et al. 2002; Macedo and Cardoso 2005]. In human beings, exploration is often driven by curiosity, a motivating force for attention focus, determination of interest, and gathering knowledge. Based on these observations, researchers devised artificial forms of curiosity for agent decision making in unknown environment. For example, Macedo and Cardoso [2001; 2002; 2005] modeled human forms of surprise and curiosity in case-based reasoning frameworks to guide agent's exploration in unknown environments populated with objects. Graziano et al. [2011] also discussed the application of computational curiosity to solve autonomous exploration problems.

To summarize, computational curiosity has the potential to contribute to various aspects of robotic systems, including selective attention and skill development. Numerous studies on computational curiosity are continuously emerging in this field [Stojanov et al. 2006; Macedo 2010; Oudeyer 2012].

5.3. Artificial Creativity

Computational creativity explores the possibility of machines capable of generating creative artifacts that are commonly defined as being previously unknown, useful and surprising [Boden 2009]. Computational curiosity has been studied in machines to demonstrate creative behaviors.

Based on Csikzentmihalyi's system view of creativity, Saunders [2007] postulated two levels of creativity: the individual level and the society level. According to Saunders, there are two questions to be answered at the individual level of creativity: (1) how to evaluate creativity, and (2) how to produce creativity. Saunders and Gero [2004] argued that curiosity can be used to guide problem solving by finding interesting design solutions as well as discovering interesting design problems. Curious design agents were proposed to evaluate the interestingness (creativity) of designs based on novelty. Research of curiosity at the society level looks into the socio-cultural influence on creativity and the acceptance of creative works by other individuals. Based on previous works, Saunders [2011] studied the society level of creativity by creating a virtual society populated with curious design agents. Simulation results showed that the artificial society exhibited certain similar behaviors as in real human societies.

Schmidhuber [2006; 2007; 2009b] explored the relationship between creativity, artists, humor and fun. He argued that art and creativity can be seen as the by-products of curiosity rewards. He further argued that the optimal curious reward framework can be sufficiently formal and precise to allow the implementation on computers and developmental robots. Schmidhuber [2006] generalized his simple principle of curiosity to form artificial creativity: "the current compressor of a given subjective observer tries to compress his history of acoustic and other inputs where possible" and "the compression progress becomes the wow-effect or intrinsic reward for the 'creative' action selector, which is thus motivated to create more data that allows for more wow-effect". Later, Schmidhuber [2011] proposed a greedy but practical implementation of the basic principle of creativity: POWERPLAY, which can automatically invent or discover problems to train a general problem solver from scratch. In his survey, Schmidhuber [2009a] drew a comparison between his formal theory with less formal works in aesthetics theory and psychology.

5.4. Games

With the advances in computer technologies for graphics, processing power, and networking, virtual worlds are emerging as platforms for massive online games. Merrick and Maher [2009] highlighted the need for new non-player characters to cope with the increasing complexity and functionality of multi-user virtual worlds. They argued that the behaviorial traits of humans and animals generated by curiosity can also advance the performance of artificial agents when dealing with complex or dynamic environments, where only limited information is available and the information changes over time.

To cope with the difficulty of predefining task specific rules or environment-specific motivation signals, Merrick et al. [2008] introduced a mechanism of curiosity into non-player characters, which enables them to direct attention to relevant information and be curious about changes in the environment. Simulation results showed that the curious motive led the non-player characters to demonstrate higher variety and complexity in behavior patterns [Maher et al. 2008].

While most of the works that apply computational intelligence to games focused on the generation of behaviors, strategies or environment, Togelius and Schmidhuber [2008] looked at the very heart of games: the rules that define a game. They proposed automatic game designs based on Koster's theory of fun and Schmidhuber's theory of artificial curiosity, where Schmidhuber's theory of curiosity is a coarse approximation of Koster's theory of fun [Koster 2005], i.e., a game is fun if it is learnable but not trivial. Li and Riedl [2010] proposed that plotlines in role-playing games can be adapted to elicit curiosity in order to improve the replay value of games, but they did not specify concrete mechanisms for modeling curiosity.

5.5. Artificial Intelligence in Education

Artificial Intelligence enables the design of pedagogical agents that can play multiple roles such as tutors, tutee or companions. These pedagogical agents are important ingredients to form social learning environments in computerized teaching and learning environment.

Pedagogical agents as tutors have been designed for achieving instructional functions [Johnson 2001]. Feedbacks from students play an important role in the instruction process to help virtual tutors adapt instructional strategies. However, most of pedagogical agents are designed to collect static feedbacks at a particular time with prescribed questions. If a pedagogical agent is curious during the instruction process, then it may provide more wits by asking situational and personalized questions, that can let students notice the weakness and conflicts in their own knowledge structures, which may induce students' curiosity.

Another role pedagogical agents often undertake is tutee. Teachable agent is a special type of tutee agent [Blair et al. 2007]. The pedagogical meaning of such an agent is that "when students teach, they tend to learn deeper". However, existing teachable agents are mostly passive learners who absorb whatever students feed them, which renders these agents less responsive during the collaborative learning process. Curiosity can make teachable agents become inquisitive in what the student teach and selective in topics based on their learnt capability.

In educational studies, it has been shown that curiosity is an important motive that drives self-direction and interests in learning. Stumm et al. [2011] attributed curiosity as the third pillar of academic performance, in addition to intelligence and conscientiousness. However, how to utilize computerized pedagogical techniques to stimulate students' curiosity is still open for discussion. Wu et al. [2012] have done some piolet studies in this direction, in which they simulated curious learning companions that can make use of users' knowledge and environmental knowledge to detect potentially interesting learning objects for users. Later, they simulated a curious peer learner in virtual learning environments to mimic the behavior of a natural student [Wu and Miao 2013].

5.6. Affective Computing

Affective computing is computing that relates to, arises from, or deliberately influences emotion or other affective phenomena [Picard 1997]. It requires multidisciplinary knowledge such as psychology, cognitive science, computer science, and engineering. Affective computing is gaining rapid popularity and has great potential in the next generation of human-computer interfaces. Curiosity is closely related to emotional constructs such as "fear", "pleasure", "boredom", and "anxiety". Computational curiosity offers a new dimension from which emotions can be appraised, apart from the consequences of events, the actions of agent, and the characteristics of objects [Ortony et al. 1988]. The consideration of computational curiosity in affective modeling is especially interesting in learning contexts and social contexts, where curiosity-related emotions significantly influences the believability and performance of emotional agents.

5.7. Artificial Companions

Artificial companions, designed to develop a close and long-term human computer relationship, have emerged in the latter half of the 2000s. Two key words, *close* and *long-term*, have been guiding the development in this field. Researchers are working on the design of believable human-computer interfaces to provide close interactions (e.g. embodied conversational agent) and robust memory architectures to sustain long-term relationships [Bickmore and Picard 2005; Wilks 2010]. Computational curiosity can be an important dimension to be studied in artificial companions for enhancing both the closeness of interactions and the possibility for long-term relationships. The potential for computational curiosity in creating a closer human-computer relationship draws evidence from psychological findings that curiosity plays an important role in promoting the intimacy of interpersonal relationships in social context. A curious artificial companion can be more responsive; may infuse more novel twists of excitement into interactions, and might induce a natural flow of engagement between the interaction discourses. As for promoting long-term relationships, curiosity can be a motivational force to learn more about the partner and develop a historical knowledge base through interactions. A curious artificial companion may be more interested to know the partner; may be more inquisitive to novel changes of the partner; and may incorporate information of the partner into part of the cognitive development of the companion itself.

5.8. Persuasive Technology

Persuasive technology deals with the use of computing systems, devices or applications to gradually change a person's attitudes or behavior [Fogg 2002]. This technology has the potential to bring constructive changes in health science, safety and education. Examples include a digital doll to persuade kids to eat fruit and vegetables, and a virtual coach to persuade the elderly to exercise more [Fogg 1999]. Understanding users' curiosity and properly infusing curiosity stimuli into the human-computer interaction process can potentially help intelligent agents achieve persuasive goals. For example, if a sales agent can successfully elicit the customer's curiosity in a product, there will be a higher chance for this product to be sold. Curiosity has been harnessed to "persuade" programmers to increase the correctness in end-user programming [Wilson et al. 2003].

5.9. Agent Negotiation

Negotiation is a process that involves two or more parties to reach an agreement. This mechanism has received increasing attention in multi-agent systems for managing inter-agent dependencies in real time [Jennings and Wooldridge 2002]. Traditional implementations of negotiation process focused on its rational aspects to build consensus. Recently, Broekens et al. [2010] argued that negotiation is a multifaceted process in which affect plays an important role. Computational curiosity has the potential to influence the human-agent negotiation process by promoting positive emotional states, responsiveness and engagement. Enabling a negotiation agent to understand the curiosity exhibited by a user may allow it to notice the unusual, surprising, or conflicting information offered by the user and reach agreements that are more socially optimal. A negotiation agent that can adapt its decision-making based on the users' curiosity may improve its chance of gaining more utility out of the final agreement.

5.10. Trustworthy Computing

Another important issue in multi-agent systems is trust management. It is useful in open and dynamic systems such as peer-to-peer systems, semantic Web, ad hoc networks, and e-commerce, etc. [Ramchurn et al. 2004; Yu et al. 2010]. Similar to negotiation, trust management is also closely related to emotion states [Dunn and Schweitzer 2005; Schoorman et al. 2007]. The motivational role of curiosity in building interpersonal relationships can contribute to the trust building between strangers [Kashdan et al. 2011]. Computational curiosity can potentially enhance an agent's judgement by making the agent more sensitive to novel, surprising, conflicting, and uncertain information presented in the environment.

ACKNOWLEDGMENTS

This research is supported in part by Interactive and Digital Media Programme Office (IDMPO), National Research Foundation (NRF) hosted at Media Development Authority (MDA) of Singapore under Grant No.: MDA/IDM/2012/8/8-2 VOL 01.

REFERENCES

G. Baldassarre. 2011. What are intrinsic motivations? A biological perspective. *In: Proc. of IEEE Conf. on Development and Learning* (2011), 1–8.

A.G. Barto, S. Singh, and N. Chentanez. 2004. Intrinsically motivated learning of hierarchical collections of skills. *In: Proc. of Intl. Conf. on Development Learn* (2004), 112–119.

D.E. Berlyne. 1960. *Conflict, arousal, and curiosity*. McGraw-Hill New York.

T.W. Bickmore and R.W. Picard. 2005. Establishing and maintaining long-term human-computer relationships. *ACM Trans. on Computer-Human Interaction* 12, 2 (2005), 293–327.

K. Blair, D. Schwartz, G. Biswas, and K. Leelawong. 2007. Pedagogical agents for learning by teaching: Teachable Agents. *Educational Technology* 47, 1 (2007), 56 – 61.

M. A. Boden. 2009. Computer Models of Creativity. *AI Magazine* 30, 3 (2009), 23–34.

J. Broekens, C.M. Jonker, and J.J.C. Meyer. 2010. Affective negotiation support systems. *Journal of Ambient Intelligence and Smart Environments* 2, 2 (2010), 121–144.

W. Burgard, M. Moors, and F. Schneider. 2002. Collaborative exploration of unknown environments with teams of mobile robots. *Advances in plan-based control of robotic agents* 2466 (2002), 187–215.

D.A. Cohn, Z. Ghahramani, and M.I. Jordan. 1996. Active Learning with Statistical Models. *Journal of Artificial Intelligence Research* 4 (1996), 129–145.

J.R. Dunn and M.E. Schweitzer. 2005. Feeling and believing: the influence of emotion on trust. *Journal of personality and social psychology* 88, 5 (2005), 736–748.

V. Fedorov. 1972. *Theory of optimal experiments*. Academic Press.

B.J. Fogg. 1999. Persuasive technologie. *Commun. ACM* 42, 5 (1999), 26–29.

B.J. Fogg. 2002. Persuasive technology: using computers to change what we think and do. *Ubiquity* 2002, 5 (2002), 27–29.

V. Graziano, T. Glasmachers, T. Schaul, L. Pape, G. Cuccu, J. Leitner, and J. Schmidhuber. 2011. Artificial Curiosity for Autonomous Space Exploration. *Acta Futura* 4 (2011), 41–51.

D.O. Hebb. 1949. *The organization of behavior*. New York: Wiley.

W. James. 1950. *Principles of psychology*. New York: Holt.

N.R. Jennings and M.J. Wooldridge. 2002. *Agent technology: foundations, applications, and markets*. Springer.

W.L. Johnson. 2001. Pedagogical agent research at CARTE. *AI Magazine* 22, 4 (2001), 85–94.

C. Karaoguz, E. Drix, D. anad Potapova, and M. Huelse. 2011. Curiosity driven exploration of sensory-motor mappings. *Deliverable for the IM-CLeVeR Spring School at the Capo Caccia Cognitive Neuromorphic Engineering Workshop* (2011), 1–7.

T.B. Kashdan. 2010. *Curious?: Discover the missing ingredient to a fulfilling life*. Harper Paperbacks.

T.B. Kashdan, P.E. McKnight, F.D. Fincham, and P. Rose. 2011. When curiosity breeds intimacy taking advantage of intimacy opportunities and transforming boring conversations. *Journal of Personality* 79, 6 (2011), 1369–1402.

T.B. Kashdan, P. Rose, and F.D. Fincham. 2004. Curiosity and exploration: Facilitating positive subjective experiences and personal growth opportunities. *Journal of Personality Assessment* 82, 3 (2004), 291–305.

R. Koster. 2005. *A theory of fun for game design*. Paraglyph Press.

Boyang Li and Mark O. Riedl. 2010. An Offline Planning Approach to Game Plotline Adaptation. *In: Proc. of the 6th Conference on Artificial Intelligence for Interactive Digital Entertainment* (2010), 45–50.

M. Li and P.M.B. Vitanyi. 2008. *An introduction to Kolmogorov complexity and its applications*. Springer-Verlag New York Inc.

J.A. Litman and C.D. Spielberger. 2003. Measuring epistemic curiosity and its diversive and specific components. *Journal of Personality Assessment* 80, 1 (2003), 75–86.

G. Loewenstein. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin* 116, 1 (1994), 75–98.

L. Macedo. 2010. Selecting Information based on Artificial Forms of Selective Attention. *In: Proc. of European Conference on Artificial Intelligence* (2010), 1053–1054.

L. Macedo and A. Cardoso. 1999. Towards artificial forms of surprise and curiosity. *In: Proc. of the European Conf. on Cognitive Science* (1999), 139–144.

L. Macedo and A. Cardoso. 2001. Modeling forms of surprise in an artificial agent. *In: Proc. of the Annaul Conf. of the Cognitive Science Society* (2001), 588–593.

L. Macedo and A. Cardoso. 2002. Assessing creativity: the importance of unexpected novelty. *In: Proc. of the ECAI Workshop on Creative Systems: Approaches to Creativity in AI and Cognitive Science* (2002), 31–37.

L. Macedo and A. Cardoso. 2004. Exploration of unknown environments with motivational agents. *In: Proc. of the Intl. Joint Conf. on Autonomous Agents and Multiagent Systems* (2004), 328–335.

L. Macedo and A. Cardoso. 2005. The role of surprise, curiosity and hunger on exploration of unknown environments populated with entities. *In: Proc. of Portuguese Conf. on Artificial Intelligence* (2005), 47–53.

M.L. Maher, K. Merrick, and R. Saunders. 2008. Achieving creative behaviour using curious learning agents. *In: Proc. of AAAI Spring Symposium on Creative Intelligent Systems* (2008), 26–28.

K. Merrick. 2008. Designing toys that come alive: curious robots for creative play. *In: Proc. of Intl. Conf. on Entertainment Computing* (2008), 149–154.

K. Merrick and E. Huntington. 2008. Attention focus in curious, reconfigurable robots. *In: Proc. of Australian Conf. on Robotics and Automation* (2008).

K. Merrick and M.L. Maher. 2009. *Motivated reinforcement learning: curious characters for multiuser games*. Springer.

K. Merrick, M.L. Maher, and R. Saunders. 2008. Achieving adaptable behaviour in intelligent rooms using curious supervised learning agents. *In: Proc. of Conf. on Computer Aided Architectural Design Research in Asia* (2008), 185–192.

S.J. Moorehead, R. Simmons, and W.L. Whittaker. 2001. Autonomous exploration using multiple sources of information. *In: Proc. of IEEE Intl. Conf. on Robotics and Automation* (2001), 3098–3103.

H. Ngo, M. Luciw, A. Forster, and J. Schmidhuber. 2012. Learning skills from play: Artificial curiosity on a Katana robot arm. *In Proc: Intl. Conf. on Neural Networks* (2012), 1–8.

M. Ogino, M. Kikuchi, and M. Asada. 2006. How can humanoid acquire lexicon?-Active approach by attention and learning biases based on curiosity. *In: Proc. of IEEE Intl. Conf. on Intelligent Robots and Systems* (2006), 3480–3485.

A. Ortony, G. Clore, and A. Collins. 1988. *The cognitive structures of emotions*. Cambridge University Press.

P.Y. Oudeyer. to appear in 2012. Developmental Robotics. *Encyclopedia of the Sciences of Learning* (to appear in 2012).

P.Y. Oudeyer and F. Kaplan. 2004. Intelligent adaptive curiosity: a source of self-development. *In: Proc. of Intl. Workshop on Epigenetic Robotics* (2004), 127–130.

P.Y. Oudeyer and F. Kaplan. 2007. What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurorobotics* 1, 6 (2007), 257–262.

P.Y. Oudeyer, F. Kaplan, V.V. Hafner, and A. Whyte. 2005. The playground experiment: Task-independent development of a curious robot. *In: Proc. of the AAAI Spring Symposium on Developmental Robotics* (2005), 42–47.

S. Pang, S. Ozawa, and N. Kasabov. 2009. Curiosity driven incremental LDA agent active learning. *In: Proc. of Intl. Joint Conf. on Neural Networks* (2009), 2401–2408.

L. Pape, C.M. Oddo, M. Controzzi, C. Cipriani, A. Frster, M.C. Carrozza, and J. Schmidhuber. 2012. Learning tactile skills through curious exploration. *Frontiers in Neurorobotics* 6 (2012), 6.

R. Picard. 1997. *Affective Computing*. The MIT Press.

S.D. Ramchurn, D. Huynh, and N.R. Jennings. 2004. Trust in multi-agent systems. *The Knowledge Engineering Review* 19, 1 (2004), 1–25.

R. Saunders. 2002. *Curious design agents and artificial creativity: A synthetic approach to the study of creative behaviour*. Ph.D. Dissertation. University of Sydney.

R. Saunders. 2007. Towards a computational model of creative societies using curious design agent. *In: Proc. of Intl. Conf. on Engineering Societies in the Agents World VII* (2007), 340–353.

R. Saunders. 2011. Artificial creative systems and the evoluation of language. *In: Proc. of the Intl. Conf. on Computational Creativity* (2011), 36–41.

R. Saunders and J.S. Gero. 2001. A curious design agent. *In: Proc. of Conf. on Computer Aided Architectural Design Research in Asia* (2001), 345–350.

R. Saunders and J.S. Gero. 2004. Curious agents and situated design evaluations. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 18, 2 (2004), 153 – 161.

J. Schmidhuber. 1991a. *Adaptive confidence and adaptive curiosity*. Technical Report FKI-149-91. Technische Universitat Munchen.

J. Schmidhuber. 1991b. Curious model-building control systems. *In: Proc. of IEEE Intl. Joint Conf. on Neural Networks* (1991), 1458–1463.

J. Schmidhuber. 1991c. A possibility for implementing curiosity and boredom in model-building neural controllers. *In: Proc. of the Intl. Conf. on Simulation of Adaptive Behavior: From Animals to Animats* (1991), 222–227.

J. Schmidhuber. 1999. Artificial curiosity based on discovering novel algorithmic predictability through coevolution. *In: Proc. of the Congress on Evolutionary Computation* (1999), 1612–1618.

J. Schmidhuber. 2002. Exploring the Predictable. *Advances in Evolutionary Computing* (2002), 579–612.

J. Schmidhuber. 2006. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science* 18, 2 (2006), 173–187.

J Schmidhuber. 2007. Simple algorithmic principles of discovery, subjective beauty, selective attention, curiosity & creativity. *In Proc: Intl. Conf. on Discovery Science* (2007), 26–28.

J. Schmidhuber. 2009a. Formal theory of creativity, fun and intrinsic motivation. *IEEE Trans. on Autonomous Mental Development* 2, 3 (2009), 230–247.

J. Schmidhuber. 2009b. Simple algorithmic theory of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. *Journal of the Society of Instrument and Control Engineers* 48, 1 (2009), 21–32.

J. Schmidhuber. 2011. Powerplay: Training an increasingly general problem solver by continually searching for the simplest still unsolvable problem. *arXiv preprint arXiv:1112.5309* (2011).

F.F. Schmitt and R. Lahroodi. 2008. The epistemic value of curiosity. *Educational Theory* 58, 2 (2008), 125–148.

F.D. Schoorman, R.C. Mayer, and J.H. Davis. 2007. An integrative model of organizational trust: Past, present, and future. *Academy of Management review* 32, 2 (2007), 344–354.

P.D. Scott and S. Markovitch. 1989. Learning novel domains through curiosity and conjecture. *In: Proc. of Intl. Joint Conf. on Artificial Intelligence* (1989), 669–674.

S. Singh, A.G. Barto, and N. Chentanez. 2004. Intrinsically motivated reinforcement learning. *In: Proc. of Intl. Conf. on Neural Information Processing Systems* (2004).

C.D. Spielberger and L.M. Starr. 1994. *Curiosity and exploratory behavior.* NJ: Lawrence Erlbaum Associates, 221–243.

G. Stojanov, A. Kulakov, and D. Clauzel. 2006. On curiosity in intelligent robotic systems. *In: Proc. of AAAI Fall Symposium on Interaction and Emergent Phenomena in Societies of Agents, Arlington, Virginia* (2006), 44–51.

J. Storck, S. Hochreiter, and J. Schmidhuber. 1995. Reinforcement driven information acquisition in non-deterministic environments. *In: Proc. of the Int. Conf. on Artificial Neural Networks* (1995), 159–164.

S.V. Stumm, B. Hell, and T. Chamorro-Premuzic. 2011. The hungry mind. *Perspective on Psychological Science* 6, 6 (2011), 574–588.

J. Togelius and J. Schmidhuber. 2008. An experiment in automatic game design. *IEEE Symposium On Computational Intelligence and Games* (2008), 111–118.

E. Uğur, M.R. Dogar, M. Cakmak, and E. Sahin. 2007. Curiosity-driven learning of traversability affordance on a mobile robot. *In: Proc. of IEEE Intl. Conf. on Development and Learning* (2007), 13–18.

R.W. White. 1959. Motivation reconsidered: The concept of competence. *Psychological Review* 66, 5 (1959), 297–333.

Y. Wilks. 2010. *Close engagements with artificial companions: Key social, psychological, ethical and design issues*. John Benjamins Publishing Company.

A. Wilson, M. Burnett, L. Beckwith, O. Granatir, L. Casburn, C. Cook, M. Durham, and G. Rothermel. 2003. Harnessing curiosity to increase correctness in end-user programming. *In: Proc. of SIGCHI Conf. on Human Factors in Computing Systems* (2003), 305–312.

Q. Wu and C. Miao. 2013. Modeling curiosity-related emotions for virtual peer learners. *Computational Intelligence Magazine* (2013), Accept.

Q. Wu, C. Miao, and Z. Shen. 2012. A curious learning companion in virtual learning environment. *In: Proc. of IEEE Intl. Conf. on Fuzzy Systems* (2012), 1–8.

W.M. Wundt. 1874. *Grundzüde physiologischen psychologie*. W.Engelman.

H. Yu, Z. Shen, C. Miao, C. Leung, and D. Niyato. 2010. A survey of trust and reputation management systems in wireless communications. *Proc. IEEE* 98, 10 (2010), 1755–1772.

Received Aug 2012; revised Jan 2013; accepted Feb 2013