On Curiosity in Intelligent Robotic Systems

Georgi Stojanov and Andrea Kulakov*

American University of Paris
France

*Faculty of Electrical Engineering
University of Sts Cyril and Methodius in Skopje
Macedonia

gstojanov@aup.fr, kulak@etf.ukim.edu.mk

Abstract

Curiosity is a notion that is increasingly receiving special attention, particularly in the context of the emerging fields of developmental robotics. In the first part of the paper we give a brief critical overview of the research in motivational systems in intelligent robotics. The overall impression is that the prevailing understanding of curiosity is rather one dimensional and reductionist in spirit. We argue that this is a result of rather simple agent's representations of the environment that are usually adopted. In the second part of the paper we put forward some arguments towards modeling of curiosity in context of other cognitive phenomena like *feeling of understanding*, analogy-making and expectations, integrated in a more general cognitive architecture.

Introduction

With the advent of developmental (epigenetic) robotics we are witnessing an increased interest in *motivational subsystems* for autonomous agents and especially in the notion of *curiosity.* This is not surprising as this community is particularly interested in agents that, during their development, exhibit ever more complex behavior via the much sought after *open-ended learning*. Sometimes this type of learning is also referred to as *task-independent* or *task non-specific*. Curiosity then, in this context, is understood to be the mechanism that would drive these systems to do *something* rather than *nothing*. Also, sometimes it is explicitly pointed to be a *goal-generation mechanism*.

But what exactly is *curiosity*? Often it is referred as a *drive* whose satisfaction should generate *positive emotions* in the agent. This terminology, most probably, originates in the behaviorist psychology where *curiosity drive* was introduced to account for situations where (mainly) animals were learning (to negotiate a maze, for example) even in the absence of any external reward (e.g. food to satisfy the hunger-drive). In a classical paper Hebb (1955) gives an overview of different understandings (and their origins) of the notion of drive. As Hebb writes, the apparent inability to reduce different types of behavior as being elicited by the basic drives (or motivations): hunger, pain, maternal, and sex drive, have led the researchers to introduce *curiosity drive* which would explain "both investigatory and manipulatory activities on the one hand, and exploratory, on the other" (Hebb, 1955). Examples include Berlyne (1950, 1954), Thomson and Solomon (1954) who explored the curiosity drive, and Harlow (1953) who investigated other types of motivations. Received reinforcements were categorized as *primary* or *secondary* (e.g. recognition by the society for one's actions). As it is not our purpose here to give a historical account of the development of the understanding of the curiosity, in the last paragraph of this section, we will briefly summarize what we believe is today's understanding.

Curiosity is usually related to notions like: *novelty, anticipation, surprise, epistemic-hunger, exploratory behavior, interest, play* and even *mild fear* as *a positive attraction of risk taking* and *mild frustration* during *problem solving* (Hebb, 1955). It is certainly widely considered as a precursor for intelligent behavior. Dennett in (1996) says that "[C]uriosity -epistemic hunger - must drive any powerful learning system". It is long considered that monkeys, dolphins and other mammals exhibit genuine curiosity. Research in other animal behavior like sharks' (e.g. Klimley and Ainley, 1998), shows compelling evidence that sharks exhibit some sort of curiosity and that their behavior is more than just *instinctive* (which is yet another term).

Often (Arbib and Fellous, 2004) curiosity is treated as a kind of *emotion*. There exists a vast body of research in emotions in AI. Scheutz in (2002) gives a critical overview of it. First, he presents an overview of various, often incompatible, characterizations of emotions in psychology. Scheutz then goes on to claim that

> it should not come as a surprise that this terminological and conceptual plurality too is spread throughout the AI literature.(*ibid.*)

Arbib and Fellous in (2004) put forward an understanding of emotions in their *functional context* i.e. striping them of their first person qualitative experiential attributes. They distinguish two aspects in emotions:

1) Emotional expression for communication and social coordination; and
2) Emotion for organization of behavior (action selection, attention and learning);

where the first is called *external aspect* of emotions, and the second *internal aspect.* An illustration for the first aspect may be the research of Breazeal (e.g. Breazeal, 2002).

In what follows, we will review some works that indeed have *functional* understanding of the curiosity and are concerned with its *internal aspect*.

Curiosity and Motivational Systems in Artificial Agents

In this section we provide a brief overview of prior work in motivational systems in intelligent robotics, and especially in the notion of *curiosity*. Needless to say, this is far from an exhaustive overview, but we believe that it gives the general understanding of the motivational issues by the researchers in this community.

Schmidhuber (1991) introduces the notion of curiosity in an otherwise Reinforcement Learning (RL) setup. In his agent there are 2 recurrent neural networks (RNN). The first one models the environment (by implementing a predictor mechanism: Situation1-Action-Situation2) and the second one (the *controller*) actually controls agent's behavior (i.e. chooses the next action to be executed). There are also *pain units* which get activated whenever the agent bumps into an obstacle. Positive reinforcement is given whenever the obstacles are avoided. The long term goals of the agent are to maximize the cumulative positive reinforcement and to minimize the cumulative pain. Both RNNs are updated during agent's life in the environment, using algorithm which is functionally equivalent to BPTT (Back Propagation Through Time). Certainly, the better the model RNN is the better the performance of the controller. Schmidhuber makes the distinction between *goal-directed* learning vs *explorative* learning. The latter is used "to increase the knowledge about the world". In this way, any goal-directed activity, which depends on the world model, will be indirectly improved. Therefore, for Schmidhuber, curiosity is the "explorative aspect of learning". The way this is implemented is quite straightforward: a *curiosity* reinforcement unit is added to the controller, and it gets activated proportionally to the *mismatch* between the output of the model network (i.e. the prediction) and the actually perceived reality (the current percept). The net effect of the mechanism is that the agent (in order to maximize the reinforcement) goes in the parts of its environment which are so far unknown, providing more learning opportunity for the model network. Moreover, the mechanism is implementable for *any* type of model based learning. However, a problem arises with monotonic increasing reinforcement for the curiosity unit. On a long term it may push the agent towards parts of the environment that are inherently unpredictable (white noise) or probabilistic. Thus, although Schmidhuber uses linear *mismatch-to-curiosity reinforcement* functions, in the same paper he says that far more realistic functions would take into account that high reinforcement should be given only when the predictor generates *near-misses*, discouraging high prediction errors. In a recent paper (Schmidhuber, 2006) summarizes his previous research into one basic *principle*:

> Generate curiosity reward for the adaptive action selector (or controller) in response to predictor improvements. (*ibid.*)

In our opinion a fair description would be to say that generally Schmidhuber equates curiosity with *surprise.*

Kaplan and Oudeyer (2002) report an "architecture of a self developing device" by which an AIBO robot can learn to trace successfully the brightest point (the *sensory input S*) in the image coming from its CCD camera (which the robot can move in the pan and tilt directions: the *motor output M*). There are three essential "processes" that interact with each other: *motivation*, *prediction*, and *actuation*. Motivation process is based on three motivational variables: *predictability* (how good is the *prediction* process in guessing the next S(t) given the previous (SM(t-1)), *familiarity* (how many times the robot has actually experienced that particular transition SM(t-1) to S(t)), and *stability* (how close remains S(t) to its average value). The reward function is such that the robot gets positive reinforcement if it *maximizes* stability motivational variable, and when it maximizes the first derivative of the *predictability* and *familiarity* motivational variables. This pushes the robot towards unknown situations Ss from where it can go back to known Ss and get some reinforcement. Apparently this reward policy is a variation of Schmidshuber's principle. Kaplan and Oudeyer also relate their motivational variables to the notions of *novelty* and *curiosity* as used by (Huang and Weng, 2002) and (Kulakov and Stojanov, 2002). In their recent work these researchers report of interesting emergent behavior in groups of AIBO robots controlled by variants of the above explained architecture (Oudeyer and Kaplan, 2006).

Blank et al. in (2005) identify three essential notions for autonomous development in robots: abstraction, anticipation, and self-motivation. The self-motivation subsystem

> […] indicates to the system how "comfortable" it is in the given environment. If it is too comfortable, it

> becomes bored, and takes measures to move the robot into more interesting areas. Conversely, if the environment is chaotic, it becomes over-excited and attempts to return to more stable and well known areas.

They present the initial results on a simulated agent that solves the navigational problem. For the abstraction and the anticipation subsystem Blank et al. use Self Organizing Maps (SOM) and RNNs. The implementation of the self-motivation is only hinted.

Working within the context of a research program called *autonomous mental development* (Weng et al., 2001; Weng, 2002) **Weng and Huang** (e.g. Huang and Weng, 2002) have implemented a motivational system in a physical robot called SAIL, which rewards the robot for going into *novel* situations. A novel situation is defined in terms of how different are the current sensory inputs from the ones encountered in the past. This pushes the robot towards regions where its *predictor* makes biggest errors in guessing the next sensory input, and, as expected the robot indeed improves its performance in environments that are deterministic and learnable. The problem arises in probabilistic and/or noisy environments where the robot apparently behaves randomly in order to maximize the prediction error (and the reinforcement with that).

In their recent research **Barto et al.** (e.g. Barto at al. 2004; Stout et al, 2005) generalize their traditional reinforcement learning approach (e.g. Sutton and Barto, 1998) by distinguishing between *external reinforcement* (usually given by a teacher or critic) and *internal reinforcement*. The internal reinforcement allows for *intrinsically motivated learning* which would enable the agent to learn

> *a broad set of reusable skills* for controlling its environment… [The intrinsic reward system] favors the development of broad competence rather than being directed to more specific externally-directed goals. But these skills act as the "building blocks" out of which an agent can form solutions to specific problems that arise over its lifetime. Instead of facing each new challenge by trying to create a solution out of low-level primitives, it can focus on combining and adjusting higher-level skills, greatly increasing the efficiency of learning to solve new problems. (Barto et al. 2004)

As for the implementation, their simulated agent has built-in notions of a *saliency* of particular stimuli (e.g. changes in light and sound). Therefore, whenever as a result of the agent's actions a salient stimulus is perceived, intrinsic reward is generated. The agent then uses standard RL techniques to actually learn how to produce that particular stimulus and, when it can do this reliably the intrinsic reward vanishes (i.e. the agent gets "bored"). These authors quote research from psychology (e.g. White, 1959) regarding motivations, and neuroscience (e.g. Dayan and Balleine, 2002; Kakade and Dayan, 2002) regarding the role of dopamine (a neuromodulator) in the intrinsic motivational control of animal behaviors associated with novelty and exploration.

In his recent book, **Trajkovski** (2006) describes an emergent curiosity in multi-agent societies, where agents model own 'confidence levels' in the predictions of the outcomes of their own actions in the environment. During agents' interactions among each other, they share their knowledge structures in a way that the less confident agent adopts the structures of the more confident agent.

As a **general discussion** from the above presented work we can say that in the majority of cases, the researchers have adopted a sort of common sense understanding of curiosity, which is often given in dictionaries, together with the usual anthropomorphism that goes along. Often it is treated as a motivational drive, and as a way for the agent to get additional internal reinforcement, in a manner very similar to behaviorist psychology in the '50s. At the same time, with some exceptions, in the research in artificial curiosity one can rarely find explicit references to relevant research from psychology. At best, the work of White (1959) and Piaget would be quoted. We want to argue that this is due to the fact that the RL paradigm (again, the preferred behaviorist learning paradigm) is most commonly adopted, and that in the RL framework such understanding is natural. In the classical RL framework the agent is represented via its *elementary actions*, which are applied in different *environmental states* using a *policy* which would *maximize* the *reward* (in terms of internal primary reinforcement) given by an *external teacher (critic)* when the agent arrives at a *goal-state*. During its life, the agent tries to learn a *model* of the environment (*if I am in situation S1 and apply action A1 I will end up in situation S2*). If there's no specific goal, via the introduction of curiosity the agent gets rewards whenever it *steps into the unknown*, which would hopefully improve its world model and its performance on subsequent tasks. Variation of this policy include: rewards proportional to the prediction error (when in S1 the world model predicts that the agent will get in S2 if A1 is performed; instead, the agent ends up in S3 and the error is defined as the *difference between* S2 and S3); rewards proportional to the *decrease* of this error in two consecutive trials; rewards that, except the prediction error, include other factors describing the agent history (familiarity of a particular S; probability to be in particular S, and the like). Rarely some more complex *representation* is used (i.e. something above the percepts S) although one can find abstractions like: *behaviors*, *visual know-how* (Kaplan and Oudeyer, 2003), *skills* and *options* (Precup 2000; Sutton, Precup, & Singh 1999). Certainly, this way of understanding and implementing curiosity can sometimes lead to exciting emergent behaviors (e.g. Oudeyer and Kaplan, 2006) but the overall impression is that that the prevailing understanding of curiosity is rather one dimensional, and rather reductionist in spirit. As mentioned in the introduction, curiosity is almost exclusively understood as something that animate otherwise *inert* agents. We want to put forward a thesis that without some more refined model for *internal representation* and *thinking* we cannot go much further towards higher developmental levels. One can even argue that the current framework doesn't offer ways to implement more sophisticated models of curiosity.

Elsewhere (Kulakov and Stojanov, 2002; Stojanov et al, *in press*) we argued for internal representations, and *cognitive models* in general, based on Bickhard's interactivism (Bickhard, 1993, 1998). The agents that were proposed in our previous works don't do something *only* because the curiosity drive is active. Their way of being in the world is by acting on it. In fact, they *enact* the world they are in. As Bickhard writes in (2000):

> A classical construal of motivation has been as that which induces a system to do something rather than nothing. The organism is assumed to be inertness motivated to do something […] That is, the organism is assumed to be inert unless some sort of "energy" is provided to move it. But organisms are alive, and living because they cannot stop, cannot be inert, without simply ceasing to exist as living beings. Living beings *cannot* do nothing. So the problem of motivation cannot be that of what makes an organism something rather than nothing. The problem of motivation must be what makes an organism do one thing rather than another.

In the remaining part of the paper we will describe our understanding of curiosity and related motivational phenomena by way of presenting our agent architectures based mainly on Bickhard's interactivism,

Interactivist Representations as a Basis for Curiosity

The purpose of our architectures Petitagé (Stojanov, 1997), (Stojanov et al, 1997), (Stojanov, 2001) and Vygovorotsky (Kulakov, 1998), (Kulakov and Stojanov, 2002) was to demonstrate the process of autonomous internalization of the environment. The key notion in these architectures is the *schema mechanism.* Inspired by the Piagetian notion of *scheme* (Piaget, 1954) a schema is an ordered sequence of elementary actions that the agent is capable of performing (e.g. ABBCDCEDEE may be an example of such a schema, where A, B, C, D, E are the elementary actions). The agent has a collection of inborn schemas that are *self-motivated* to get executed. One can think of this set of schemas as a result of the evolution that has equipped the agent with useful set of initial behaviors for the survival (self-preservation) of the agent. The sensory-motor influx thus generated represents the background for the activity of *pattern detectors* whose role is to build *chunks* (proto-categories, proto-concepts) and anticipate future sensory-motor influxes. Initially, this process is guided by the primitive internal value systems based on the satisfaction of agent's drives, one of them being the curiosity drive that will be explained later. The result of this process is a graph-like internal representation of the environment like in Figure 1.

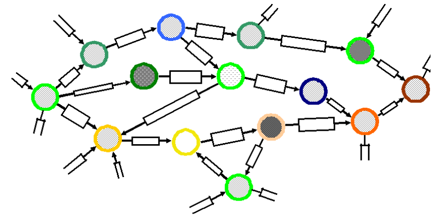

Figure 1. Graph like internal representation as a result of agent's interaction with the environment

Initially the links represent sequences of *concrete* actions and the nodes are configurations of perceptual states. The *abstraction building* mechanism introduces new links and nodes using various techniques. The new links and nodes are more and more abstract as they depart from the initial schemas grounded in the actual sensory-motor experiences. Elements of this representation also include mechanisms for *novelty detection*, judging the *reliability* or the *confidence* of the schemas' predictions, and enable development of more sophisticated models of different cognitive phenomena necessary for the development of curiosity.

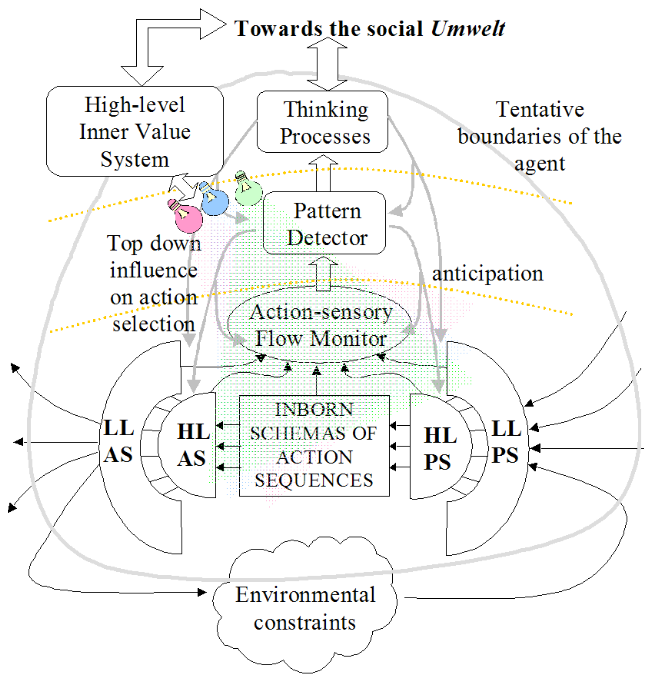

Figure 2. General cognitive architecture of an agent. **LLPS** stands for Low-Level Perceptual System, **HLPS** is High-Level Perceptual System, **LLAS** is Low-Level Action System and **HLAS** is High-Level Action System.

In Figure 2 we give a schematic representation of our generic agent architecture. The light bulbs represent the Low-Level Inner Value System, which provides different emotional 'coloring' (influence) on the whole architecture.

We can metaphorically depict the agent's initial point of view as in Figure 3a. It is intended to show the sensory-motor influx as an un-interpreted, un-segmented environment that the newborn agent finds itself in.

During the agent-environment interaction, the environment provides *stimuli* which serve as *triggers* for certain motor-schemas of the agent. The environment also imposes certain constraints on the way that the schemas are executed.

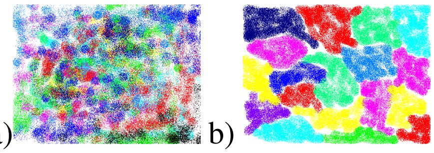

Figure 3. a) Uninterpreted, unsegmented sensory stimuli, versus b) segmented, classified sensory stimuli

Together with the biases that come from the inborn value system, this process results in emergence of certain structure in the sensory-motor influx. This is illustrated in Figure 3b. We can think of these clusters as *proto-concepts* that will eventually lead to concepts which the agent then uses to interpret (parse) the immediate environment. This is illustrated with the downward arrows in Figure 2 (*anticipation*, and *top-down influence on action selection*). Another interpretation would be that the agent is actually changing the environment it is in by actively imposing a structure on the incoming sensory influx.

We summarize the 4 crucial mechanisms that guide agent's development while it matures:

- abstraction mechanism that provides *chunking* of the sensory inputs and of the motor outputs, i.e. of the *pattern sequences (sensory-motor schemas)* and which enables the agent to deal with more and more complex situations with the same or less cognitive effort;
- thinking and planning mechanism which hypothetically *combines* various previous experiences into new knowledge, for example by *analogy-making*
- mechanism that provides emergence of more complex inner value and motivational systems according to which new experiences are judged, foreseen and executed;
- socialization mechanism that enables the agent to interpret in a special way inputs coming from other *intelligent (live or A-live) agents*.

The wide arrows at the top of Figure 2, towards and from the social Umwelt are depicted here in order to stress the interaction between the agent and its social environment, although this interaction anyhow goes through the perceptual system and through the motor action system of the agent.

Figure 4 depicts our notion of a knowledge graph (internal representation) of the cognitive agent, which is implemented in our cognitive architectures. Stemming from several levels of abstraction, proto-concepts are grouped into more or less coherent and compact groups of sensory-motor pattern sequences, connected among each other with abstract action-types. Although not depicted there, action-types are also following the abstraction hierarchy from more primitive motor actions. We do not discuss here about concepts, because according to our stance, in agreement with Hofstadter's theory of Fluid Concepts (Hofstadter, 1995) and Rosch's theory of Radial Categories (Rosch, 1973), a concept is a 'hallo' of several proto-concepts and more or less abstract sensory-motor pattern sequences, activated around a certain central part of that particular concept.

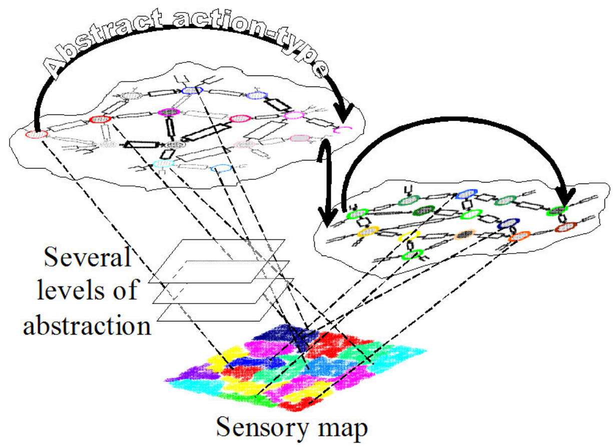

Figure 4. Knowledge structures developed in a grown-up cognitive agent. After several levels of abstraction, proto-concepts at the Sensory map are chunked into more abstract proto-concepts. These proto-concepts are grouped into more or less coherent and compact groups of sensory-motor pattern sequences connected among each other with abstract action-types.

So far, research in AI and related disciplines have not come out with an agent that will also follow this type of ontogenetic development. Researchers concentrated either on very high level representation formalisms (predicate-like, close to natural language) detached from the actual embodiment of the agent, or, disregarding any representation, they concentrate their effort on producing agents with purely reactive behavior.

In our opinion, having a curiosity drive which would only maximize the learning curve of an agent, equipped with a mechanism for reinforcement learning, is not enough to fully explain the curiosity. First of all, curiosity, being a part of the motivational system, may only partially influence the decision for taking actions. Also, the curiosity may provoke internal interest for thinking about certain parts of knowledge graph, without overt behavior, depending on the complex High-level Inner Value System, as well.

Curiosity and the feeling of understanding

In what follows we would relate the curiosity with the feeling of understanding. According to Natika Newton (Newton, 1996), we have an experience of understanding about which it can be said that, if we feel that we are perplexed, then we are, and if we feel that we understand, then no matter that it sometimes can be misunderstanding, we do feel that we understand. It is not about 'correct' meaning, it is about *meaningfulness*. Also there seem to be levels of confidence of our understanding from totally perplexed through tip-of-the-tongue, when there seems to be some initial understanding, to the *A-ha Erlebniss* when we suddenly receive an understanding of something.

We distinguish two kinds of feelings of understanding in our architecture. One is the feeling of understanding for the working memory, the other one is for the whole knowledge graph.

For each conceptual node an average reliability (or confidence) of all schemas in the vicinity of certain conceptual node (within a distance of few links) is calculated. Whenever this node becomes activated by the mechanism for activation spreading, this local average confidence is used to judge the situation at hand. Whenever this parameter is high enough, the situation at hand (represented in the Working Memory) seems to be well understood. Whenever it is low, the situation at hand is perplexing. Unlike this, the feeling of understanding for the whole knowledge graph is calculated as an average reliability (or confidence) of all the schemas in the graph.

The curiosity drive in our architectures is defined as directly proportional function of the both feelings of understanding as defined before. Its purpose is to create a tendency to raise the confidence of the agents' knowledge. The agent's interest for a certain part of the knowledge graph (or the desirability of that part) is directly proportionate to the overall average confidence of the schemas in the graph, while on the other hand, it is inversely proportionate to its local average confidence. This interest for a certain part of the knowledge graph is twofold, both for behavior, as a goal for a plan, and for thinking, as a target for analogy-making.

This is modeled so, because whenever we are generally perplexed we are hardly eager to learn about new things, but rather we first try to explain for ourselves what is perplexing us at hand. So, only when the agent is confident enough in its currently active knowledge, it is willing to continue to explore new situations or only to think about new things. This prevents the appearance of many 'holes' of loosely understood parts in agent's knowledge graph.

Having defined all this, we can mathematically model the *Aha-Erlebniss* as a first-order derivative of the feeling of understanding of the working memory.

Curiosity and analogy-making

The true understanding of each concept in the memory will be achieved when that conceptual node will be strongly connected with other nodes, not only with hierarchical links obtained through the process of abstraction, but also with reliable (confident) sensory-motor schema links to other conceptual nodes as well.

> One understands an object if one can imagine it incorporated in a token of an understood action-type with an image rich enough to guide oneself in reaching the goal. (Newton, 1996)

From this we can conclude that it is not necessary to have immediate experience for all of our understanding. It is sufficient that we have rich enough abstract experience which we can relate by *analogy* to the current interest.

Back in terms of our architectures, the percept (or the group of already recognized proto-concepts) will be understood if a connection can be found between the current percept and some percepts that are part of some already understood (reliable, confident) sensory-motor schema. A transfer of knowledge occurs when a good analogy-mapping has been made. New relations (schemas) are constructed between the nodes of the 'target situation' and they are given a particular initial reliability (confidence) depending on the 'goodness' of mapping. Since in the moments before this transfer occurs, the 'target situation' was loosely connected with the rest of the knowledge graph, adding new connections would increase the interconnectivity of the whole graph. In that way the feeling of understanding will be increased, i.e. a moment of understanding has happened, no matter that these new relations can sometimes lead to a misunderstanding.

The understanding of one part of conceptual network (the target) metaphorically by another part (the source), is actually a transfer of the relations (schemas) that hold in the source part to the target part.

Curiosity and expectations

Our architectures have one very important characteristic of intelligence – the expectations. The mind is fundamentally an anticipator, as Dennett deduced succinctly in his efforts for explaining intelligence (Dennett, 1996). In every moment, in every step we anticipate something, and then we expect the outcome of our actions. By learning the environment in a manner of connecting the current percept with some motor scheme to some expectation for the next percept and so on, our agents structure the environment according to their possibilities to perceive, to act, and to sense. In that way they build an expectation framework for the environment which they inhabit. So, whenever the expectations are met, the agent does not have to bother what it will do next – it continues the sensory-motor schema as if on "auto-pilot". The problem for the agent appears when the agent is surprised by the detected mismatch between the expectations and the current percepts. Only when it is surprised, the agent has to stop and try to figure out the solution according to the current state in the stack of goals. Besides this, new goals have to be redeemed when the current plan is finished, if the stack of goals happens to be empty.

Still these surprises can be internally provoked, as for example when a new way of understanding something that we knew before comes to our mind or a new plan has just been constructed. We would like to distinguish between these kinds of surprises generated by internal processes with the surprises generated by the mismatch during the process of behavior monitoring.

> For the behavior it is inherent to reduce uncertainty, while for the processes of thinking or imagination it is inherent to produce unexpectancies [hypotheses] that will later be checked. (Kulakov, 1998)

While the curiosity drive has a function for reducing the uncertainty of the knowledge by proving it or rejecting it by behavior, it also has a function to increase the unexpectancy or uncertainty, by adding new nodes in the knowledge graph during *imagination* or *thinking*, done by analogy-making for example. This is obtained by calculating the desirability of each part of the knowledge graph according to this drive and choosing some of them, not only as goals for behavior plans, but also as goals (targets) for making analogy mappings and transfers.

Summary and Conclusions

We have opened this paper by reviewing representative current research on curiosity in artificial agents. The overall impression is that the prevailing understanding of curiosity is rather one dimensional and reductionist in spirit. We argued that this was a result of rather simple agent's representations of the environment that were usually adopted. In the second part of the paper we briefly summarized our generic cognitive architecture based on Bickhard's interactivism and Piaget's notion of scheme. We argued that a cognitive architecture built along these principles would allow for more complex internal representation, which in turn would allow more sophisticated models of curiosity. Curiosity would not be treated as a simple driving force which only pushes the agent to do something, but rather as an elaborated mechanism which is inseparable from the internal knowledge representation and can guide the processes of thinking and imagination. The anticipations, feeling of understanding, and analogy making are complementary processes that together with curiosity enrich internal representation. Ultimately, the whole process would culminate in language-like representation which will open the way to powerful symbolic reasoning available only to linguistically competent agents.

Acknowledgment

The work described in this article has been partially funded by the European Commission's Sixth Framework Programme under contract no. 029427 as part of the Specific Targeted Research Project XPERO ("Robotic Learning by Experimentation").

References

Arbib, M. A. and Fellous, J-M., 2004. Emotions: from brain to robot, *TRENDS in Cognitive Sciences* Vol.8 No.12

Barto, A. G., Singh, S., and Chentanez, N., 2004. Intrinsically motivated learning of hierarchical collections of skills. In *Proceedings of International Conference on Developmental Learning (ICDL).* MIT Press, Cambridge, MA.

Berlyne, D. E., 1950. Novelty and curiosity as determinants of exploratory behavior. British Journal of Psychology, 41, pp. 68-80.

Berlyne, D. E., 1954. A theory of human curiosity. *British Journal of Psychology*, 45, pp. 180-191.

Berlyne, D. E. (1960). Conflict, Arousal, and Curiosity. New York: McGraw Hill

Bickhard, M. H., 1993. Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, pp. 285-333

Bickhard, M. H., 1998. Levels of representationality. Journal of Experimental and Theoretical Artificial Intelligence, 10(2), pp. 179-215

Bickhard, M. H., 2000. Motivation and Emotion: An Interactive Process Model. In R. D. Ellis, N. Newton (Eds.) *The Caldron of Consciousness.* J. Benjamins. pp. 161-178.

Breazeal C., 2002 *Designing Sociable Robots*, The MIT Press

Blank, D.S., Kumar, D., Meeden, L., and Marshall, J. 2005. Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture. *Cybernetics and Systems*, 36(2), Taylor & Francis, pp. 125-150

Dayan, P., and Balleine., W., 2002. Reward, motivation and reinforcement learning. *Neuron*, 36, pp. 285–298

Dennett, D., 1996. *Kinds of Minds*, Basic Books, New York.

Frijda, N. and Swagerman, J., 1987. Can computers feel? Theory and design of an emotional model. *Cognition and Emotion.* 1 (3). pp 235-357

Hofstadter, D., 2001. Analogy as the Core of Cognition, in Dedre Gentner, Keith J. Holyoak, and Boicho Kokinov (eds.) *The Analogical Mind: Perspectives from Cognitive Science,* Cambridge, MA: The MIT Press/Bradford Book, pp. 499-538

Harlow, H. F., 1953. Mice, monkeys, men, and motives. *Psychological Review*, 60, pp. 23-32

Hebb, D. O., 1955. Drives and the C.N.S. (Conceptual Nervous System), *Psychological Review*, 62, pp. 243-254.

Hsiao, K. and Roy, D., 2005. A Habit System for an Interactive Robot. *AAAI Fall Symposium 2005: From Reactive to Anticipatory Cognitive Embodied Systems*

Huang, X. and Weng, J. 2002. Novelty and reinforcement learning in the value system of developmental robots. In *Proceedings of the 2nd international workshop on Epigenetic Robotics* - Lund University Cognitive Studies 94, pp. 47-55.

Kaelbling, L. P., Littman, M. L., and Moore, A. W. 1996. Reinforcement learning:
A survey. *Journal of Artificial Intelligence Research*, 4, pp. 237-285

Klimley, A. P., Ainley, D. G., (eds.) 1998. *Great White Sharks: The Biology of Carcharodon carcharias*, Academic Press

Kakade. S. and Dayan, P., 2002. Dopamine: Generalization and bonuses. *Neural Networks*, 15, pp. 549–559

Kuipers, B., Beeson, P., Modayil, J., and Provost, J. 2006. Bootstrap learning of foundational representations, *Connection Science,* 18 (2), pp 145-158

Kulakov, A., 1998. Vygovorotsky - A Model of an Anticipative and Analogy-making Actor. MSc Thesis, New Bulgarian University, Sofia, Bulgaria

Kulakov, A. and Stojanov, G., 2002. Structures, inner values, hierarchies and stages: Essentials for developmental robot architectures. In *Proceedings of the 2nd International workshop on Epigenetic Robotics* - Lund University Cognitive Studies 94, pp. 63-69

Newton, N. 1996. "Foundations of Understanding", John Benjamins Publishing Company, Amsterdam.

Oudeyer, P.-Y. and Kaplan, F. 2006. The discovery of communication. *Connection Science*, 18(2).

Piaget, J. 1954. The Construction of Reality in the Child, Basic, New York

Precup, D. 2000. *Temporal Abstraction in Reinforcement Learnin.* Ph.D. Dissertation, U. of Massachusetts, Amherst.

Scheutz, M., 2002. Agents With or Without Emotions? In *Proceedings of FLAIRS 02*, AAAI Press, pp. 89-94

Schmidhuber, J., 1991. A possibility for implementing curiosity and boredom in model-building neural controllers. In J. A. Meyer and S. W. Wilson, editors, *Proceedings of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, MIT Press/Bradford Books, pp. 222-227

Sloman, A. 1987. Motives mechanisms and emotions. Emotion and Cognition, 1(3):217-234. Reprinted in M. A. Boden, editor, *The Philosophy of Artificial Intelligence*, Oxford Readings in Philosophy Series, Oxford University Press, pages 231-247, 1990.

Stojanov, G., 1997. Expectancy Theory and Interpretation of EXG Curves in Context of Biological and Machine Intelligence. Unpublished PhD Thesis, ETF, Sts. Cyril and Methodius University, Skopje

Stojanov, G., Bozinovski, S., Trajkovski G., 1997. Interactionist-Expectative View on Agency and Learning. *IMACS Journal for Mathematics and Computers in Simulation,* North-Holland Publishers, Amsterdam.

Stojanov, G., 2001. Petitagé: A case study in developmental robotics, *Proceedings of the 1st International Workshop on Epigenetic Robotics*, Lund University Cognitive Studies, 85

Stojanov, G., Trajkovski, G., Kulakov, A., *in press,* Interactivism in AI and Intelligent Robotics, *New Ideas in Psychology*, Elsevier

Sutton, R. S. and Barto, A. G., 1998. *Reinforcement Learning*, The MIT Press.

Sutton, R., Precup, D. and Singh, S., 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning, *Artificial Intelligence* 112(1-2), pp.181–211.

Thompson, W. R., & Solomon, L. M., 1954. Spontaneous pattern discrimination in the rat. *J. comp. physiol. Psychol.,* 47, 104-107.

Trajkovski, G. 2006. An Imitation-based approach to modeling homogenous agents societies, Idea Group Publishing, Hershey.

Velásquez, J., 1998. When Robots Weep: Emotional Memories and Decision Making. In: *Proceedings of AAAI-98.*

Weng, J., McClelland, J., Pentland, A. Sporns, O., Stockman, I., Sur, M., and Thelen, E. 2001. Autonomous mental development by robots and animals. *Science* 291, pp. 599–60.

Weng, J., 2002. A theory for mentally developing robots. In *Second International Conference on Development and Learning*. IEEE Computer Society Press

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