The role of Surprise, Curiosity and Hunger on Exploration of Unknown Environments Populated with Entities

Luís Macedo, Amílcar Cardoso

***Abstract*— This paper describes an approach based on affect to the problem of exploring unknown environments populated with entities by agents. To this end, a multi-agent system based on the notion of affect as well as on the Belief-Desire-Intention (BDI) model was used. The affective component of the agents is confined to the motivations that are usually associated to exploratory behavior: surprise, curiosity and hunger. An experiment that evaluates the role of these motivations in exploration performance is presented.**

***Index Terms*—Exploration of unknown environments, Curiosity, Hunger, Surprise.**

I. Introduction

Exploration gathers information about the unknown. Exploration of unknown environments by artificial agents (usually mobile robots) has actually been an active research field. The exploration domains include planetary exploration (e.g., Mars or lunar exploration), search for meteorites in Antarctica, volcano exploration, map-building of interiors, etc. The main advantage of using artificial agents in those domains instead of humans is that most of them are extreme environments making exploration a dangerous task for human agents. However, there is still much to be done especially in dynamic environments as those real environments mentioned above. Those real environments consist of objects. For example, office environments possess chairs, doors, garbage cans, etc., cities comprise several kinds of buildings (houses, offices, hospitals, churches, etc.), cars, etc. Many of these objects are non-stationary, that is, their locations may change over time. This observation motivates research on a new generation of mapping algorithms, which represent environments as collections of objects. Moreover, the autonomy of agents still needs to be improved, as happens for instance in planetary exploration which is still too human dependent. Several exploration techniques have been proposed and tested either in simulated and real, indoor and outdoor environments, using single or multiple agents (for an overview see e.g., [1-3]). In human beings, exploration has been closely connected with motivation (including emotion and drives). This relationship between exploration and motivation has been defended for a long time in the realms of psychology and ethology. James' concept of selective attention [4], Freud's term cathexis [5], and McDougall's notion of curiosity instinct [6] are foundation thoughts for the relationship between motivation and exploratory behaviour. Therefore, a reasonable approach is to model artificial agent's exploration according to humans, i.e., in a human-like fashion by assigning artificial agents mentalistic qualities such as emotion and motivation, beliefs, intentions, and desires. Actually, there is one primary reason for taking the way humans explore the environment as a reference: the problem of modelling exploration in humans has already been successfully solved by millions of years of evolution. Yet, in general, a lot of barriers have been found to incorporate models of emotion in artificial agents. Research in AI has almost ignored this significant role of emotions on reasoning, and only recently this issue was taken seriously (e.g., [7-18]) mainly because of the recent advances in neuroscience, which have given evidence that cognitive tasks of humans, and particularly planning and decision-making, are influenced by emotion [19].

In this paper we describe an approach based on affect to the problem of exploring unknown environments by agents. We developed a multi-agent system based on the notion of affect as well as on the BDI model, which was used as a platform to develop the application to exploration of unknown environments with affective agents. Primary relevance is given to the architecture of an affective agent and especially to its affective module and its influence on exploratory behavior. We confined the set of motivations to those that are more related with exploratory behavior in humans [20].

The next section describes the approach for exploring unknown environments with affective agents. Section 3 presents an experiment that was conducted to evaluate that approach. Finally, we present conclusions.

This work was supported in part by the PRODEP.

L. Macedo is with the Instituto Superior de Engenharia de Coimbra, Quinta da Nora, 3030 Coimbra Portugal (e-mail: lmacedo@isec.pt) and with the Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Polo II, 3000 Coimbra Portugal (e-mail: macedo@dei.uc.pt).

A. Cardoso is with the Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Polo II, 3000 Coimbra Portugal (e-mail: amilcar@dei.uc.pt).

II. Exploring Unknown Environments with Affective Agents

We study the problem of exploring unknown environments using a multi-agent system based on the notion of affect as well as on the BDI model. This multi-agent system is suitable to applications in which the entities (agents) are distributed in a physical environment. This is the case of the domain of exploration of unknown environments. The simulation environment considered as a test bed to our approach to exploration comprises therefore a variety of entities located at specific positions. In this case, the objects are confined to buildings. The *structure* of the buildings comprises the shape (triangular, rectangular, etc.) of the roof, facade, door and windows. The possible *functions* may be: house, church, hotel, hospital, etc.

The architecture that we adopted for an agent (Figure 1) is based on the BDI approach [21]. As in many other agents' architectures, the architecture followed in our work includes the following modules: sensors/perception; effectors/actuators; memory/beliefs; emotions, drives and other motivations (or simply motivation); intentions/goals; desires; and, deliberative reasoning/decision-making. The deliberative reasoning/decision-making module is in the core of the architecture. It receives internal information (from memory) and environment information (through the sensors) and outputs an action that has been selected for execution. The process of action selection takes into account the states of the environment the agent would like to happen (desires), i.e., it selects an action that leads to those states of the environment the agent prefers. This preference is implicitly represented in a mathematical function that evaluates states of the world in terms of the positive and negative feelings they elicit in the agent. Thus, this function obeys to the Maximum Expected Utility principle [22]. In this case, the utility is in positive feelings. The intensities of these feelings (motivations) are computed by the motivation module taking into account both the past experience (the information stored in memory) and the present environment description provided by the sensors.

To explore the environment, each agent is continuously performing the deliberative reasoning/decision-making algorithm. Thus, each agent at a given time senses the environment to look for entities and compute the current world state (location, structure and function of those entities) based on the sensorial information and on the generation of expectations for the missing information. Then, a goal of kind *visitEntity* is generated for each unvisited entity (including those within the visual range and also those out of this range that were previously percepted but not yet visited). In addition, a goal of the kind *visitLoc* is generated for all the frontier cells [3]. Then, these goals are then inserted in the ranked list of goals which might already contain previous goals generated in the past but not yet accomplished. This list of goals is ranked according to the Expected Utility (EU) of the goals, which is computed based on the intensities of motivations predicted as explained below.

The next three subsections describe in more detail the main modules of the architecture.

A. *Agent's Memory*

The memory of an agent stores information about the world. This information includes the configuration of the surrounding world such as the position of the entities (objects and other animated agents) that inhabit it, the description of these entities themselves, descriptions of the sequences of actions (plans) executed by those entities and resulting from their interaction, and, in generally, beliefs about the world. This information is stored in several memory components. Thus, there is a (grid-based) metric map [23] to spatially model the surrounding physical environment of the agent. Descriptions of entities (physical structure and function) (Figures 2 and 3) and plans are stored both in the episodic memory and in the semantic memory [24]. The physical structure of an entity may be described analogically or propositionally [25]. The function is simply a description of the role or category of the entity in the environment. For instance, a house, a car, a tree, etc. Like the description of the physical structure, this may be probabilistic because of the incompleteness of perception. Concrete entities (i.e., entities represented in the episodic memory) with similar features may be generalized or abstracted into a single one, an abstract entity, which is stored in the *semantic memory for entities*. Figure 3 presents a semantic memory obtained from the episodic memory of entities shown in Figure 2.

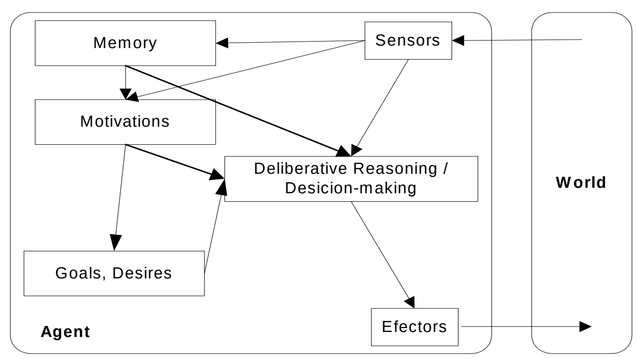

Fig. 1. Agent's architecture.

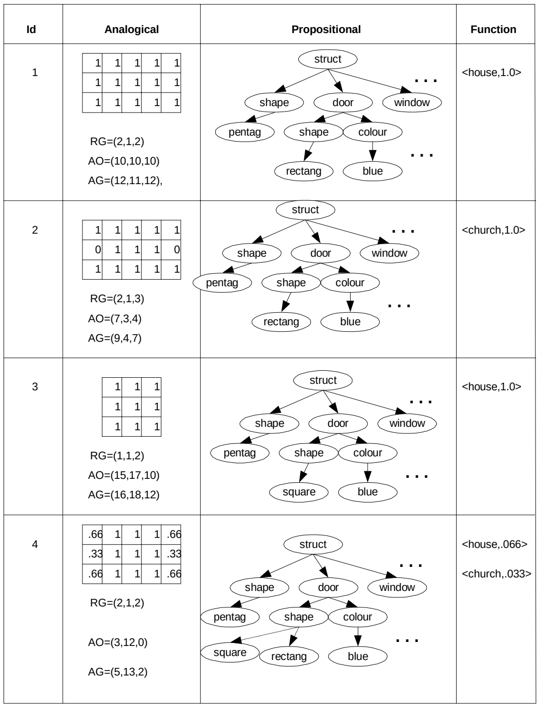

Fig. 2. Episodic memory of entities.

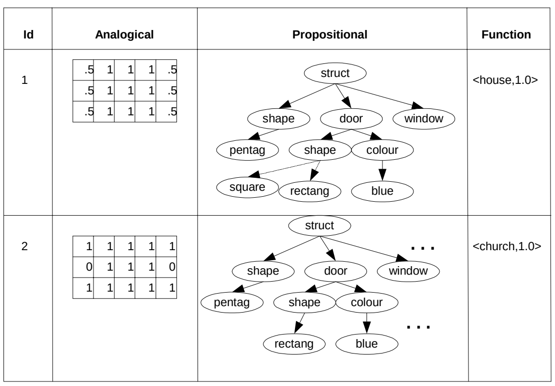

Fig. 3. Semantic memory of entities.

B. Motivations

The module of Motivations receives information from the current state of the environment and outputs the intensities of emotions, drives, and other motivations. In this paper, this module is confined to the motivations that are related with variables that directly instigate exploration: surprise (elicited by unexpectedness), curiosity (elicited by novelty and uncertainty), and hunger (the drive that reflects the need of an energy source).

1) Surprise

A reasonable approach to model the agent's surprise function is according to that of humans. Experimental evidence from human participants summarized in [26] suggests that the intensity of felt surprise increases monotonically, and is closely correlated with the degree of unexpectedness (see [27] for more details). This means that unexpectedness is the proximate cognitive appraisal cause of the surprise experience. On the basis of this evidence, it is reasonable that the surprise "felt" by an agent elicited by an event X is proportional to the degree of unexpectedness of X. The intensity of surprise elicited by X should therefore be an (at least weakly) monotonically increasing function of 1-P(X) [27]. However, an additional empirical and theoretical study [28] conducted in the domains of political elections and sport games with several surprise functions suggests that the surprise felt by an agent elicited by an event Eg, g ∈ {1, 2, …, m}, among a set of m mutually exclusive events E={E1, E2, …, Em} is given by:

$$
\begin{aligned}
&SURPRISE(E_g) = UNEXPECTEDNESS(E_g) = \\
&= \log_2(1 + P(E_h) - P(E_g))
\end{aligned}
\qquad (1)
$$

In this formula, E_h, $h \in \{1, 2, \ldots, m\}$, is the event with the *highest* probability of the set *E*. It implies that, within each set of mutually exclusive events, there is always at least one (E_h) whose occurrence is entirely unsurprising, namely the event with the maximum probability in the set (P(E_h)). For the other events X in the set, the surprise intensity caused by their occurrence is the logarithm of the difference between P(E_h) and their probability P(E_g) plus 1. This difference can be interpreted as the amount by which P(E_g) has to be increased for E_g to become unsurprising. This equation predicts that that maximum surprise, i.e., SURPRISE(E_g) = 1, occurs only if P(E_h) = 1 and hence, by implication, P(E_g) = 0. (In the Ortony and Partridge model [29], this corresponds to those situations where the disconfirmed event E_h is immutable, i.e., its probability is 1). Therefore, this formula seems to correctly describe surprise in the election example. Confirming this impression, this formula also acknowledges the intuition that if there are only two alternative events E_g and E_h (= not E_g), it predicts that E_g should be unsurprising for P(E_g) ≥ 0.5, for in this case E_g is also the event with the highest probability in the set. By contrast, for P(E_g) < 0.5, it predicts that E_g should be surprising and increasingly so the more P(E_g) approaches 0, with maximum possible surprise (SURPRISE(E_g) = 1) being experienced for P(E_g) = 0. In addition, however, it also captures the nonlinearity of the surprise function suggested by the experiments with humans reported in [26].

The above equation just gives the surprise of an event after its occurrence. However, it is possible to compute beforehand the surprise the agent expects to feel from a scenario *S* whose outcome is one the events of the set of mutually exclusive events $E=\{E_1, E_2, \ldots, E_m\}$. This is given by the following equation that resembles the equation of EU [22] as well as the equation of entropy [30], where the logarithmic factor plays the role of utility and *surprisal*[1], respectively:

$$
E[SURPRISE(S)] = \sum_{i=1}^{m} P(E_i) \times \log_2(1 + P(E_h) - P(E_i)) \qquad (2)
$$

[1] Notice that the notion of *surprisal* that belongs to information theory differs from our notion of surprise because the former does not capture correctly the human surprise

The computation of the intensity of surprise elicited by an object relies on considering the object as consisting of pieces: the cells of the analogical description, the propositions of the propositional description, and the function. Surprise is computed based on all those pieces of an object. Each piece of an object is considered as a scenario. For some of those scenarios there is already an outcome event and for others don't, but rather a set of possible events associated with a probability of occurrence. This means for the former scenarios the probability distribution contains a single pair <event, probability> – the certain event, while for the latter scenarios the probability distribution contains multiple pairs. In this case, these pairs constitute the active expectations [29] of the agent which may conflict with the information further acquired for these uncertain scenarios. Although, the probability distribution of the scenarios with no uncertainty contains a single pair, it is possible to compute the probability distribution as if there were no certainty by computing the probability for the already known events as well as for the other events that could have happen. The pairs of such probability distributions correspond to passive expectations as they are computed only after the outcome of a scenario is known. Whatever the scenario contains uncertainty or not, the probabilities of the probability distributions are computed in three manners, depending on the category of knowledge to which the piece of information belongs: (a) for the scenarios corresponding to pieces of the propositional description the Bayes' equation is used taking as evidence the rest of the pieces of the propositional description that are already known; (b) for the scenario corresponding to the function of the object, the Bayes' equation is used taking as evidence the pieces of information of the propositional description that are already known; (c) for the scenarios corresponding to the cells of the analogical description, the probabilistic analogical description is used which is obtained based on the probability distribution for the function of the object.

The intensity of surprise results from the contribution of both the pieces with no uncertainty (X_C) and the pieces with uncertainty (X_U):

$$\begin{aligned} &SURPRISE(X) = SURPRISE(X_C) + E\left[SURPRISE(X_U)\right] = \\ &= \sum_{E_g \in X_C} \log_2(1 + P(E_h) - P(E_g)) + \\ &+ \sum_{S \in X_U} \sum_{i=1}^{m_S} P(E_i) \times \log_2(1 + P(E_h) - P(E_i)) \end{aligned} \quad (3)$$

To compute the surprise for the pieces with no uncertainty, the probabilities of the event with the highest probability of the set and of the event that really occurred are retrieved from the probability distributions and used in equation 1. Then all the surprise values computed for all the pieces of the object that are known are summed. For the uncertain pieces of an object the process is similar except that all the probabilities are taken from the probability distributions and not only those of the event with the highest probability of the set and of the event that really occurred. Equation 2 is used to compute the *expected surprise values* of all the uncertain pieces and then they are summed.

2) Curiosity

We define curiosity/interest (following McDougall [6], Berlyne [20] and Shand [31]) as the desire to know or learn an object that arouses interest by being novel or uncertain, which means that novel and uncertain objects, i.e., objects with at least some parts that are not yet known, stimulate actions intended to acquire knowledge about those objects. While novelty means new information, uncertainty means that new information is probably to be acquired. Information is a decrease in uncertainty which, according to information theory, is measured by entropy [30]. An object may comprise a known part and an uncertain part. Thus, if we accept the above definition, the curiosity/interest induced in an agent by an object X depends both on the novelty or difference of X relatively to the set of objects present in the memory of the agent AgtMem, and on the entropy of the object:

$$\begin{aligned} &CURIOSITY(X) = \\ &= NOVELTY(X) + UNCERTAINTY(X) = \\ &= DIFFERENCE(X, AgtMem) + ENTROPY(X) \end{aligned} \quad (4)$$

Like in surprise, the computation of the curiosity elicited by an object is based on considering the object as consisting of pieces. Curiosity is thus computed based on all those pieces of an object. Like surprise, curiosity results from the curiosity elicited by the certain parts and the uncertain parts of the object. The pieces of the object that contain no uncertainty, i.e., which are already known, are used to compute the novelty of the object, while the uncertainty pieces[2] are used to compute the entropy of the object.

Let us consider first the certain pieces. In order to compute its novelty, the object is compared with every object in memory. This comparison may involve the comparison of the propositional and analogical descriptions, and the functions. Since the propositional description is represented in a graph-based way, the measure of difference relies heavily on error correcting code theory [32]: the function computes the distance between two objects represented by graphs, counting the minimal number of changes (insertions and deletions of nodes and edges) required to transform one graph into another. A similar procedure is applied to compute the difference between the analogical descriptions: the analogical descriptions of two objects are superimposed and then the cells that don't match are counted. The difference between two objects in what respect to the function is either 1 or 0, depending on they match or don't match. To compute the difference of a given object relatively to a set of objects, we apply the above procedure to each pair of objects formed by the given object and an object from the set of objects. The minimum of those differences is the difference of the given

[2] These could be all the pieces because the certain pieces have a null contribute to the overall entropy of the object.

object relatively to the given set of objects.

The entropy is computed based on all parts of an object that contain uncertainty. This includes the analogical (X_A) and propositional (X_P) descriptions of the physical structure, and the function (X_F):

$$\begin{aligned} H(X) &= H(X_A) + H(X_P) + H(X_F) = \\ &= \sum_{i=1}^{m} p^i \log_2(\frac{1}{p^i}) + \sum_{z=1}^{l} \sum_{j=1}^{r_z} p_j^z \times \log_2(\frac{1}{p_j^z}) + \\ &+ \sum_{K=1}^{n} p_k \times \log_2(\frac{1}{p_k}) + \\ &+ (1-p^i)\log_2(\frac{1}{1-p^i}) \end{aligned} \quad (5)$$

3) Hunger

The drive hunger is defined as the need of a source of energy. Given the capacity C of the storage of that source (C=1, i.e., C=100%), and L the amount of energy left ($0 \leq L \leq C$), the hunger elicited in an agent is computed as follows:

$$HUNGER = C - L \quad (6)$$

C. Motivations and Deliberative Reasoning/Decision-making

The motivational system plays an important role in the generation and ranking of goals/intentions which is performed by the deliberative/decision-making module. Actually, according to psychologists, motivations are the source of goals in several manners: these goals may be included in emotions (e.g., when an agent feels anger about something, a possible triggered goal might be fisting the entity that is on the origin of the anger), or emotions may be themselves the goals (e.g., an agent looks for states of the world that elicit certain positive emotions such as happiness or surprise). We take seriously this principle. Therefore, an agent selects actions or sequences of actions that lead to those states of the world that maximize positive feelings and minimize negative ones. For instance, an agent establishes the goal of visiting an object that seems beforehand interesting (novel, surprising) because visiting it will probably make it feel happy for acquiring information. After a set of goal tasks are generated, their EUs are computed. This is performed predicting the motivations (surprise, curiosity and hunger) that could be elicited when the effect E_j^k of a goal task *T* takes place [17, 33]:

$$\begin{aligned} EU(E_j^k) &= \alpha_1 \times U_{surprise}(E_j^k) + \alpha_2 \times U_{curiosity}(E_j^k) + \\ &+ \alpha_3 \times U_{hunger}(E_j^k) = \\ &= \alpha_1 \times Surprise(E_j^k) + \alpha_2 \times Curiosity(E_j^k) + \\ &+ \alpha_3 \times Hunger(E_j^k) \end{aligned} \quad (7)$$

The functions $Surprise(E_j^k)$, $Curiosity(E_j^k)$, and $Hunger(E_j^k)$ are replaced by the functions of surprise, curiosity, and hunger defined above and applied to the resulting state of the world when the effect E_j^k takes place. To what parts of the state of the world they are applied is determined in the definition of each action. For instance for the case of task of visiting an entity/cell they are applied to the entity or to the cell visited. In the first case, it is restricted to surprise and curiosity, while in the second case only the function hunger is used.

III. Experimental Evaluation

The main goal of this experiment is to assess the influence of surprise, curiosity and hunger on the performance of the exploration of environments populated with entities. We let an agent explore exhaustively the environment several times, each time with a different exploration strategy (eight of the strategies result from different combination of the parameters of equation 7 – see table below – while another is based on a classical exploration strategy that takes into account the distance to traverse and the amount of information expected to be acquired). We collected the number of different entities visited along the time.

A. Materials and Method

In order to compare the performance of the different strategies we consider an agent exploring 3 simulated environments, each time with a different strategy for exploration. One of those strategies, A8, is based on the distance to be traveled by the agent and the expected information gain which is defined by the entropy [34]. Eight of these strategies result from considering the combinations of the parameters of Equation 7. The possible combinations of these parameters and the correspondent strategies are presented in Table I. With strategy A0, the agent performs undirected exploration (random) [35]. With strategy A1, the agent performs directed exploration based solely on hunger. With strategy A2 it performs directed exploration based solely on curiosity. With strategy A3, the agent performs directed exploration based on curiosity/interest and hunger. With strategy A4, the agent performs directed exploration based only on surprise. With strategy A5, it performs directed exploration based surprise and hunger. With strategy A6, it performs directed exploration based on surprise and curiosity. With strategy A7, it performs directed exploration based on surprise, curiosity, and hunger.

TABLE I
Units for Magnetic Properties

Strategy	α1 - Surprise	α2 - Curiosity	α3 - Hunger
A0	0	0	0
A1	0	0	-1
A2	0	1	0
A3	0	1	-1
A4	1	0	0
A5	1	0	-1
A6	1	1	0
A7	1	1	-1

The 3 simulated environments in which the agent was ran contain an average of 60% of entities that are each other similar. Environments that contain some entities that are equal seem to be appropriate to test this matter because the agent has to select the entities to maximize the number of different entity models acquired. This is harder and clearly seen in environments with a considerable amount of equal entities than in environments in which all or almost all the entities are equal or different.

The procedure of this experiment consists simply in running the agent nine times in all the environments each time with a different strategy, starting from the same location. The visual range of the agent was also constant (10 cells). The value of the number of different entities visited was collected. This is then the dependent variable, while the strategy and the environment are the independent variables. This enables us to take conclusions about how the influence of the strategy on the value of "the different entities visited" evolves during the exploration of an environment. This means that it is possible with this experiment to see the influence at any time during exploration. The main motivation for this study is that sometimes there is a time limit to explore an environment that is too short to explore it completely. Hence, with this experiment we may take conclusions about which strategy or strategies are better for these situations.

B. Results

The results of this experiment are presented in Figure 4. It plots the time series of the variable "number of different entities visited" for the strategies A1, A2, A3, A4 and A8. For the sake of simplicity we avoid plotting the other strategies (A7 is very similar to A3, A6 to A2, A5 to A4, and A0 was clearly the worst strategy). It can be seen that the strategy A1 outperforms the others clearly after the complete exploration of the environment, although A3 and A2 are very close contenders. It is worth of notice that A3, A2 and A8 produced better results than A1 in a quite large time interval. We verify that the number of entities visited using these strategies are higher than those of using the other strategies from time 34 to time 51. Another result is that the classical strategy is outperformed by most of the strategies when there is complete exploration of the environment.

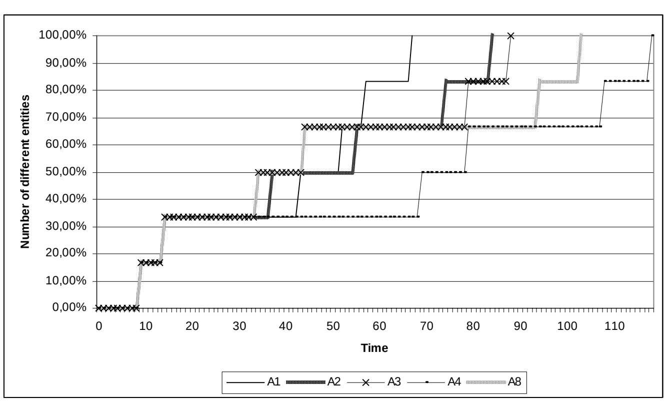

Fig. 4. Time series of the variable "number of different entities visited" for the different exploration strategies.

C. Discussion

This experiment shows that the agent explores faster all the different entities of the environment when it uses a strategy that takes into account hunger either alone or combined with surprise and/or curiosity. Actually, an agent that uses those strategies visits faster all the different entities than if it uses the other strategies. In this case an agent that uses these strategies that take into account surprise or curiosity, jointly or independently, together with hunger the erratic paths that are achieved when using surprise or curiosity alone gives rise to ordered exploration paths and hence to a significative increase on efficiency. Actually, the motivation to visit entities or frontier cells that are expected to elicit curiosity and/or surprise but that are far away from the location of the agent is restrained by the hunger that is expected to be felt on those destination locations. Curiosity and to some extent surprise lead the agent to visit entities that are expected to be different. So, it could be expected that when they are taken alone or combined the agent would visit more different entities. However, this does not happen because the agent looses most of the time traversing long distance, although this depends on the configuration of the environment. So, the role of hunger is essential to restrain this impetus to visit far and expected new entities. The classical exploration strategy performs worse mainly because it does not take into account the novelty of the entities.

When there is a short time limit to explore the environment (between 34 and 51), the experiment shows that the strategies A3 or A8 outperform the others. The worse results of strategy A1 seem to show that it is too sensitive to the distribution of the entities in the environment. Actually, it does not provide the agent with capability of looking for maximal knowledge. Instead, it enables the agent to avoid spending energy.

IV. Conclusion

We have presented an approach for directed exploration of unknown environments based on surprise, curiosity and hunger. The strategy that takes into account hunger seems to be the best strategy after exploring exhaustively an environment. However, when there is a shorter time limit to explore an environment, the strategy that takes into account curiosity and hunger seems to be better. Surprise seems to be unnecessary for exploration. However, it proved to be useful when exploration is performed in the context of creativity, i.e., when the primary goal of the agent is not to gain knowledge but instead to admire entities that are considered artistically or scientifically creative. Therefore, surprise seems to be useful in exploration tasks performed in environments such as museums.

In the future we expect to extend this experiment so that the approach could be tested in much more environments in order to achieve statistically significative results. Strategy A1 seems to be too sensitive to the distribution of the entities in the environment. Therefore, additional tests are required to take definitive conclusions. Besides, in addition to the variable "number of different entities visited" there other variables that are important to measure and compare. One of those variables is the number of measurements performed [34]. Another future work is about the computation of novelty or difference which gives equal relevance to all the features of the objects. This approach based on old notions of similarity in psychology has long since been discredited by data.

REFERENCES

[1] S. J. Moorehead, R. Simmons, and W. L. Whittaker, "Autonomous exploration using multiple sources of information," in *Proceedings of the IEEE International Conference on Robotics & Automation*, 2001.

[2] W. Burgard, M. Moors, and F. Schneider, "Collaborative exploration of unknown environments with teams of mobile robots," in *Advances in plan-based control of robotic agents*, vol. 4266, *Lectures Notes in Computer Science*, M. Beetz, J. Hertzberg, M. Ghallab, and M. Pollack, Eds. Berlin: Springer Verlag, 2002.

[3] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes, "Coordination for Multi-Robot Exploration and Mapping," in *Proceedings of the AAAI-2000*, 2000.

[4] W. James, *The principles of psychology*. New York: Holt, 1890.

[5] S. Freud, *The basic writings of Sigmund Freud*. New York: Random House, 1938.

[6] W. McDougall, *An introduction to social psychology*. London: Methuen, 1908.

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

[8] E. Oliveira and L. Sarmento, "Emotional advantage for adaptability and autonomy," in *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. New York, NY: ACM Press, 2003, pp. 305 - 312.

[9] W. Reilly, "Believable Social and Emotional Agents,". Pittsburg, PA: School of Computer Science, Carnegie Mellon University, 1996.

[10] R. Picard, *Affective Computing*. Cambridge, MA: MIT Press, 1997.

[11] C. Elliott, "The Affective Reasoner: A Process Model of Emotions in a Multi-agent System," Northwestern University, The Institute for the Learning Sciences, Ph.D. Thesis 32, 1992.

[12] M. Maçãs, R. Ventura, L. Custódio, and C. Pinto-Ferreira, "DARE: An Emotion-Based Agent Architecture," in *Proceedings of the Fourteenth International Florida Artificial Intelligence Research Society Conference*, I. Russell and J. Kolen, Eds.: AAAI Press, 2001, pp. 150-154.

[13] A. Paiva, J. Dias, D. Sobral, R. Aylett, P. Sobreperez, S. Woods, C. Zoll, and L. Hall, "Caring for Agents and Agents that Care: Building Empathic Relations with Synthetic Agents," in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, N. Jennings and M. Tambe, Eds.: ACM Press, 2004, pp. 194-201.

[14] C. Martinho, M. Gomes, and A. Paiva, "Synthetic Emotension," in *Proceedings of the 4th International Workshop on Intelligent Virtual Agents*, T. Rist, R. Aylett, and D. Ballin, Eds. Berlin: Springer-Verlag, 2003, pp. 57-61.

[15] L. Botelho and L. Coelho, "Artificial Autonomous Agents with Artificial Emotions," in *Proceedings of the second international conference on Autonomous agents*, K. Sycara and M. Wooldridge, Eds. New York: ACM Press, 1998, pp. 449-450.

[16] D. Cañamero and J. Fredslund, "How does it feel? Emotional Interaction with a Humanoid LEGO Robot," in *Socially Intelligent Agents: The Human in the Loop. Papers from the AAAI 2000 Fall Symposium*, K. Dautenhahn, Ed. Menlo Park, CA: AAAI Press, 2000, pp. 23-28.

[17] C. Castelfranchi, R. Conte, M. Miceli, and I. Poggi, "Emotions and goals," in *Perspectives on Cognitive Science*, B. Kokinov, Ed. Sofia: New Bulgarian University, 1996, pp. 131-145.

[18] J. Gratch, "Why you should buy an emotional planner," in *Proceedings of the Agents'99 Workshop on Emotion-based Agent Architectures*, 1999.

[19] A. Damásio, *Descartes'error, Emotion Reason and the Human Brain*. New York: Grosset/Putnam Books, 1994.

[20] D. Berlyne, "Novelty and curiosity as determinants of exploratory behavior," *British Journal of Psychology*, vol. 41, pp. 68-80, 1950.

[21] M. Bratman, D. Israel, and M. Pollack, "Plans and resource-bounded practical reasoning," *Computational Intelligence*, vol. 4, pp. 349–355, 1988.

[22] S. Russell and P. Norvig, *Artificial Intelligence - A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall, 1995.

[23] S. Thrun, "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel, Eds. San Mateo, CA: Morgan Kaufmann, 2002.

[24] E. Tulving, "Episodic and Semantic Memory," in *Organization of Memory*, E. Tulving and W. Donaldson, Eds. New York: Academic Press, 1972, pp. 381-403.

[25] M. Eysenck and M. Keane, *Cognitive psychology*. London: Lawrence Erlbaum Associates, 1991.

[26] R. Reisenzein, "The subjective experience of surprise," in *The message within: The role of subjective experience in social cognition and behavior*, H. Bless and J. Forgas, Eds. Philadelphia, PA: Psychology Press, 2000.

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

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

[29] A. Ortony and D. Partridge, "Surprisingness and Expectation Failure: What's the Difference?," in *Proceedings of the 10th International Joint Conference on Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann, 1987, pp. 106-8.

[30] C. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379-423;623-656, 1948.

[31] A. Shand, *The foundations of character*. London: Macmillan, 1914.

[32] R. Hamming, "Error Detecting and Error Correcting Codes," *The Bell System Technical Journal*, vol. 26, pp. 147-60, 1950.

[33] R. Reisenzein, "Emotional Action Generation," in *Processes of the molar regulation of behavior*, W. Battmann and S. Dutke, Eds. Lengerich: Pabst Science Publishers, 1996.

[34] C. Stachniss and W. Burgard, "Exploring unknown environments with mobile robots using coverage maps," in *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, G. Gottlob and T. Walsh, Eds. San Francisco, CA: Morgan Kaufmann Publishers, Inc., 2003, pp. 1127-1132.

[35] S. Thrun, "Efficient Exploration in Reinforcement Learning," Carnegie Mellon University, Computer Science Department, Pittsburgh, PA CMU-CS-92-102, 1992.