Towards an artificial user: the "what" problem for an architecture capable of developing new goals

Riccardo Manzotti, Vincenzo Tagliasco
DIST, University of Genoa,
Viale Causa 13, I-16145 Genova (Italy)

Summary

A user is a system capable of creating and pursuing individual goals. Is it possible to design and implement an artificial user? Traditional artificial systems focus on how achieving a given goal. Most learning algorithms look for an optimal solution of a problem, given a set of optimization criteria and a goal (or a set of goals). However, real agents and real users have to develop new goals in order to cope with their environment. They must find "what" they want to achieve and not only "how". The development of completely new goals on the basis of the interaction with the environment is here defined the "what" problem. In this paper we will try to define it and we will propose an architecture capable of addressing it. Such an architecture is proposed as the foundation for an artificial user.

Keywords: ontogeny, phylogeny, motivations, robot, development, goal

1. What and how

What is a user? A user is a system with its own goals actively trying to purse them by means of tools or interacting with other systems. How does a cognitive system, like that of human beings, produce its own goals? And why do current artificial systems lack this capability?

Current implementations of artificial systems focus on the implementation of intelligent algorithms to implement optimal performances to achieve a fixed goal. Human users show a different approach. They not only achieve an optimal performance with respect to a given set of hardwired goals; they develop unpredictable and unexpected new goals. The new goals are totally or partially dependent on a limited set of hardwired goals fixed at design time; they are the result of the interaction with the environment and can overcome the original ones.

The objective of this paper is twofold: to outline the "what" problem and to describe an architecture capable of addressing it.

Artificial systems are designed with an already fixed set of goals that has to be achieved. Designers focus their efforts to find "how" these goals can be achieved.

Learning is defined as a modification in our behaviour driven by a goal. Various learning paradigms focus on this modification of the behaviour: Supervised and Unsupervised Learning and Reinforcement Learning are valid examples. For instance, Sutton and Barton claim that "reinforcement learning is learning what to do – how to map situations to actions – so as to maximize a numerical reward signal" (Sutton and Barto 1998): the goal is expressed by a reward signal which is *a priori* defined. They claim that “the basic idea is simply to capture the most important aspects of the real problem facing a learning agent interacting with its environment to achieve a goal. […] All reinforcement learning agents have explicit goals. These explicit goals are fixed at design time. We could say that Reinforcement Learning deals with situation in which the agent seeks to achieve a goal despite uncertainty about its environment.

We want to address, in this paper, the process by which it is obtained “what” (the goal) the agent has to achieve. This is the “what” problem. Using the Reinforcement Learning terminology we could say that the “what” problem is equivalent to looking for new reward functions. In Reinforcement Learning systems “the reward function must necessarily be unalterable by the agent”. On the contrary many biological systems are capable of developing partially or totally unpredictable goals. We propose to design systems that develop their own reward functions: systems capable of becoming artificial users.

A few examples.

Romeo meets Juliet at a Capuleti party. Nothing in his upbringing forced him to be in love with her. However after his first encounter, to meet her again is his most important goal. All his behavior is directed to achieve this new goal.

After a few minutes a mallard duck opens its eyes and sees, as a first large moving object, the face of Konrad Lorenz. Subsequently it tries to keep it in its field of view as much as possible. It aims at this new goal.

In the aforementioned examples, a new goal is added to the causal structure embodied by an agent. Biological agents, by means of natural selection or personal experience are able to do so. .In them goals must “self emerge” since they cannot receive their goal from the outside

To develop new goals is important since the environment cannot be completely predicted at design time. In nature there are basically two situations in which new goals are added: phylogenesis and ontogenesis. During phylogenesis a new species develops completely new goals that arise from the modification in the environment that are the product of the emergence of new ecological balances. During ontogenesis, accordingly to the degree of behavioral plasticity, each individual is capable of developing new goals from its experience. New goals can be a better specification of the original one or can be completely new. For instance imprinting in mallard ducks can be seen as simple example of specification of sub goals. The more general hardwired goal of “looking for a mother” becomes the specific goal of looking for a given specimen of mallard duck (or Konrad Lorenz himself!).

2. Architectures for agents: a taxonomy

Not all the goals of biological systems are fixed at birth: they only possess a very limited, survival driven, built-in set of goals. As they grow, biological systems generate new goals on the basis of two separate factors: their genetic background and their past experience. Both are necessary to select a particular goal.

Behaviour-based artificial structures develop on experience and goals defined elsewhere at design time (Arkin 1999). Goals are the result of the interaction between experience and a limited number of hardwired instincts (the ones provided by genes). In many complex biological systems, it is possible to distinguish between phylogenetic aspects and ontogenetic ones, nature versus nurture (Gould 1977; Elman, Bates et al. 2001; Ridley 2004). Instincts are not goals. They are not the result of an ontogenetically caused motivation. They are a set of procedures to produce goals, as the imprinting procedure of birds.

What is a goal? When an agent wants an event to happen again, a new goal springs up. In this respect a goal is projected towards the future; this is only half the story. Why the agent should be willing such an event to happen again? Obviously because of something that happened in the past. In this respect a goal is a projection of the past. Goals are causal structures that link the past with the future. Because of this, they are the ideal candidate as a building block for an agent.
Is it possible to implement instincts and goals in an artificial system?
We propose a taxonomy of architectures: a fixed control architecture, a learning architecture and a goal generating architecture (Figure 1). In the first case, the system has no capability of modifying *how* it does *what* it does. There is a input-output mapping module, which takes the input signal and produces the output on the basis of some a priori hard-wired module. Examples of this structure are simple control devices or machine automata. In the second case, the system is capable of modifying its behavior to fulfill some a priori target. The system is capable of modifying *how* it behaves. Examples of this structure are reinforcement learning or supervised learning artificial neural networks. In the third case, the system is capable of modifying not only *how* it does what it does, but also to define *what* it does. The goal module sets the goals that have to be pursued by the learning module.
On the basis of these three levels, three kinds of architectures are proposed: fixed control architectures, learning architectures, goal generating architectures.

2.1 Fixed control architecture (input-output)

In this case, the causal structure of the system is fixed. There is no ontogenesis whatsoever. Notwithstanding the behavioral complexity of the system, everything happens because it has been previously coded. A mechanical device and a complex software agent are not different in this respect: both are pre-programmed in what they must achieve and how they must achieve it. Nothing in their structure is caused by their experiences. Suitable examples of this category are Tolam's artificial sow bug, Braitenberg's thinking vehicles (Braitenberg 1984), Brooks artificial insects and recent entertainment robots like Sony AIBO and Honda's humanoid ASIMO (2002).

2.2 Architecture for learning ("how")

A different level of structural dependency on the environment is provided by the architectures that can learn *how* to perform a task. Behaviour-based robots can be classified in this category. Systems based on artificial neural networks are well-known examples of this kind of architecture. These systems determine how to get a given result once they have been provided with a specific goal. The goal can be given either as a series of examples of correct behavior (supervised learning) or as a simple evaluation of the global performance of the system (reinforcement learning) (Sutton and Barto 1998). In both cases some kind of learning is applied. These systems lack the capability of creating new goals. There are several examples of this kind of learning agent: Babybot at LIRA-Lab (Metta, Manzotti et al. 2000), Cog at MIT (Brooks, Breazeal et al. 1999).

2.3 Goal generating architecture ("what")

A system that learns both *how* to perform a given task and *what* task must be performed, corresponds to an ontogenetic architecture. This is the case for most, if not all, mammals; it is true for primates and for human beings. They are systems capable of developing new goals that do not belong to their genetic background. In the field of artificial systems there has been a series of attempts to address this problem (Fukushima, Okada et al. 1994; Weng 1996) as well as attempts to locate similar structures in the cortical architecture of humans (Togawa and Otsuka 2000). For their development, these systems depend more on the environment than the previous two categories. A system belonging to the first category does not depend on the environment for what it does or for how it does what it does..

3. The what problem

An ideal autonomous system must be able to develop in a completely unknown environment. In order to do that, it must be provided with sensory and motor capabilities. It is reasonable to assume that the sensory and motor capabilities delimit a more limited world than the environment itself. The world offered by the senses is only a possible world: it contains many more events that those to which the system can react to.

We can make examples even in the case of humans. We believe in a world in which we are possible causally connected to all kinds of visual and sensorial patterns. However, for a normal user it is impossible to recognize the difference between a Pepsi and a Coca. Many people are not able to read and, as a result, they are not able to recognize printed characters.

The following model is proposed. The environment is defined as the set of all possible events in the surrounding of a given agent (World 3). The sensory and motor capabilities define a second subset of the first one. Only those events that *can* produce an effect (directly or indirectly) in the agent by means of its body structure are part of the second subset (World 2). However this does not entail that the agent is causally related with all of them. This would the same as claiming that just because someone has a pair of ears s/he should be able to have a reaction for all possible kind of music and different language. This is not the case. Only a limited number of them (World 1) have been selected during the past phylogenetic and ontogenetic history of the agent and only those constitute the relevant environment for it (Figure 2a).

A caveat is here needed. When we present these sets we are not referring to sets of objects. We refer to set of events. From the point of view of the user, s/he is causally related to events, not to objects. For instance s/he reacts to an execution of a symphony, s/he reacts to the showing of an image, s/he reacts to someone shouting. By means of the events the agent is usually able to relate with objects and other more or less static entities. However the first contact with the world is based on events, not on objects. It is arguable that objects are superset of events constructed on the basis of events (Manzotti and Tagliasco 2001; Manzotti 2003).

What do we mean here for an event to belong to the agent's world? The concept is straightforward. An event belongs to the agent's world if and only if, when it takes place, it produces a distinguishable effect in the agent. What do we mean for a distinguishable effect? A few examples are useful. If a flow of high energy x rays is projected against an unknown subject, it will produce with high probability several harmful effects on the cellular structure of his/her body. However, it would go unnoticed. If a pattern is showed to some uninterested subject it would produce effects (in the subject's retina a flow of chemical activity would surely be correlated with the visual property of the pattern), however it will have no relevance for his/her following activity. On the other hand if a new person is introduced, it will normally produce in a subject a lasting capability of reacting to his/her visual appearance. When a new word is heard, it produces a new capability in the subject. Namely, the capability of reacting to that sequence of phonemes. If “qyxtzy” is heard, it does not produce any effect in English speaking listeners. If “love” is heard, some kind of effect goes on.

What is the difference between the two? A first difference is that the couple (event and agent) must be able to produce a recognizable and repeatable effect. Is there any simple criterion to identify “a distinguishable effect”? Yes there is. It is the fact that the system must be able to use that event as a goal. This is the crucial difference. When something happens, it must not only produce an effect in the agent, but the effect it produces must be usable as a goal for the system as a whole.

Let us go back to the previous examples. Accidental effects and unnoticed patterns cannot be used as goals by the system. Therefore they are not distinguishable notwithstanding the strength of the effects produced. On the contrary a new face or a new word can be used as goal. “love” is a big goal for many, “qytxtzy” is not. However it could become. If for example tomorrow a subject would read an interesting tale about the adventures of a curious character called “qytxzy”, s/he would remember it and s/he could use it as a new goal.

This last example is important because it brings us back to one of the original problems: where are the relevant events coming from? They come from the past of the history of the agent. They come from the environment itself. But how? We just said that the agent is enclosed inside a world of relevant events (world 1).

Well, the answer is that the boundary between the agent’s world and the possible worlds is an open boundary. New events can drop from the world 2 in the world 1 by enlarging the world 1. Therefore a better metaphor would be one in which the boundary between world 1 and world 1 are not represented by a closed line but by a bubble like process of enlargement (Figure 2).

A general schema is provided in Figure 3. The world is perceived and acted upon through the motor and sensory bottlenecks. The agent perceives the world through the senses but it perceives only those events which are compatible with its perceptual categories. The matching between these perceptual categories and the current sensory input is fed into some behavior generator. In most cases of artificial agent this is what happens. However in an ideal system, the perceptual categories and the goals are the result of a process of epigenetic development. By the term “epigenetic” we mean that it depends on the contingent experience the system had during its own lifespan. There are two situations in which the generation of goals is important. First, controlling the learning of correct behaviors with respect to the developed goals. Second, controlling the learning of new perceptual categories.

An ideal agent is capable of selecting whatever event from the environment (among those detectable by means of its sensorial capabilities) and then of using it as a goal for further development. Since the goal comes from the environment and since the capability of detecting it must be a result of the will to detect it as a goal, there is some kind of circular causality going on. It is not an undesired result. It is what we should expect since it is what happens during phylogenesis and ontogenesis.

A few words can be spent on this issue. The selection of new events in the environment (we always refer to the world 2 in the following paragraphs) is a process that can be performed in three different ways: totally unconstrained, partially constrained and completely constrained.

The last case is equivalent to most of artificial robotic implementations around. There is a fixed set of goals, hardwired at design time. The learning is aimed at achieving these goals. There is no space left for enlarging the repertoire of relevant events in the environment. The world 3 is empty and it is not enlarging. (In the world 3 there are only the newly acquired or self emerged events). The world 3 is missing.

The totally unconstrained case is of course unrealistic since it would lack any possibility of control. The outcome of the agent development would be completely unpredictable. The enlargement of world 3 is left to chance and to the resources available to the agent.
The partially constrained case is the most interesting for us. It leaves some space for the enlargement of world 3 but it will focus its enlargement on some pre-defined criteria. For instance human beings are going to devote their perceptual capabilities to recognition of human faces much more than shapes of rock or leaves. There must be some phylogenetic bias to focus their development on specific kind of events.
The agent will internally implement the process of self-emergence of new perceptual categories. How this can be done will be explained in the following paragraph. What is interesting is that the system will produce two outputs: a first perceptual output which represents the current sensorial input *after* it has been somehow matched with the epigenetic perceptual structure and the goal signal. These two signals have a completely different logical and practical meaning but are both the result of the epigenetic history of the agent.

4. An architecture capable of developing new goals

The architecture we present in this paragraph is a potential candidate to endorse the epigenetic development of new goals we just described. A detailed description of it and some experimental results can be found in Manzotti, Tagliasco 2005 (Manzotti and Tagliasco 2005). There are probably other architectures with the same capacity. However what we present here must be seen as a valid candidate for the development of a completely adaptive agent (Figure 3).
The main goal of the architecture is to let events to become goals for the systems. Every real system doesn’t have necessarily to start from scratch. Some information can be embedded in the system in such a way as to permit the system to bootstrap itself and to take advantage more quickly of the environment. It’s necessary to provide the system with some kind of meta-functions: a function that in a given environment will provide a specific function. In nature, an example of this meta-function is provided by imprinting. The bird has no specific knowledge of the shape of its mother, but has some meta-functions (looking for a moving object, in a fixed window frame of time). In this way the meta-function will be able to build a new function.
Since these functions must be ready before the development we called them ‘phylogenetic’. They could have been as well been called ‘hard-wired’ functions. There are two kinds of phylogenetic functions: normal and meta.
Since the potential complexity of the environment is unlimited, it makes sense to have a mechanism in order to reduce the span of events to categorize and to explore. This mechanism corresponds to the phylogenetic functions. For instance, face recognition is triggered in human by some phylogenetic functions which is more interested to those events which are roughly similar to a face.
If an agent would have only phylogenetic functions it were an unsupervised learning or reinforcement learning system. It will have a set of goals fixed at design time: it will have a very limited capacity for development.
The epigenetic functions are equal to the phylogenetic ones. The only difference is that they are based on the categories developed during epigenesis. By using the categories developed during epigenesis, the system is partially unpredictable because, given a certain environment and certain phylogenetic functions, the agent will be able

to pick up only certain events. However, in practice, given an unconstrained environment, the development will be unconstrained as well.
A system must be able to associate a value (proportional to how much that particular event is important on the basis of its past history) to what happens: this signal will be subsequently used as a reinforcement signal to implement the development of new goals.
Finally the system will produce, for selected classes of events, a relevant signal which is the maximum among all the phylogenetic functions and the epigenetic ones. It tells how much the current input is representative of the past history of the system. At the beginning, this signal will only be the result of the phylogenetic functions, but when new values come in from the epigenetic functions, the relevant signal will include the new goals and will allow the system to develop new goals.

5. References

Arkin, R. C. (1999). Behavior-Based Robotics. Cambridge (Mass), MIT Press.
Braitenberg, V. (1984). Vehicles: Experiments in Synthetic Psychology. Cambridge (Mass), MIT Press.
Brooks, R. A., C. Breazeal, et al. (1999). The Cog Project: Building a Humanoid Robot. Computation for Metaphors, Analogy, and Agents. Nehaniv. Berlin, Springer-Verlag. 1562.
Elman, J. L., E. A. Bates, et al. (2001). Rethinking Innateness: A Connectionist Perspective on Development. Cambridge (Mass), MIT Press.
Fukushima, K., M. Okada, et al. (1994). "Neocognitron with Dual C-Cell Layers." Neural Networks 7(1): 41-47.
Gould, S. J. (1977). Ontogeny and Phylogeny. Cambridge (Mass), Harvard University Press.
Manzotti, R. (2003). A process based architecture for an artificial conscious being. Process theories. J. Seibt, Kluwer Academic Press. Process Theories: Crossdisciplinary studies in dynamic categories: 285-312.
Manzotti, R. and V. Tagliasco (2001). Coscienza e Realtà. Una teoria della coscienza per costruttori e studiosi di menti e cervelli. Bologna, Il Mulino.
Manzotti, R. and V. Tagliasco (2005). "From "behaviour-based" robots to "motivations-based" robots." Robotics and Autonomous Systems In press.
Metta, G., R. Manzotti, et al. (2000). Development: is it the right way towards humanoid robotics? ISA-6, Venezia, IOS Press.
Ridley, M. (2004). Nature via nurture. Great Britain, Harper.
Sutton, R. S. and A. G. Barto (1998). Reinforcement Learning. Cambridge (Mass), MIT Press.
Togawa, T. and K. Otsuka (2000). "A model for Cortical Neural Network Structure." Biocybernetics and Biomedical Engineering 20(3): 5-20.
Weng, J. J. (1996). Cresceptron and SHOSLIF: Toward comprehensive visual learning. Early visual learning. S. K. Nayar and T. Poggio. New York, Oxford University Press.

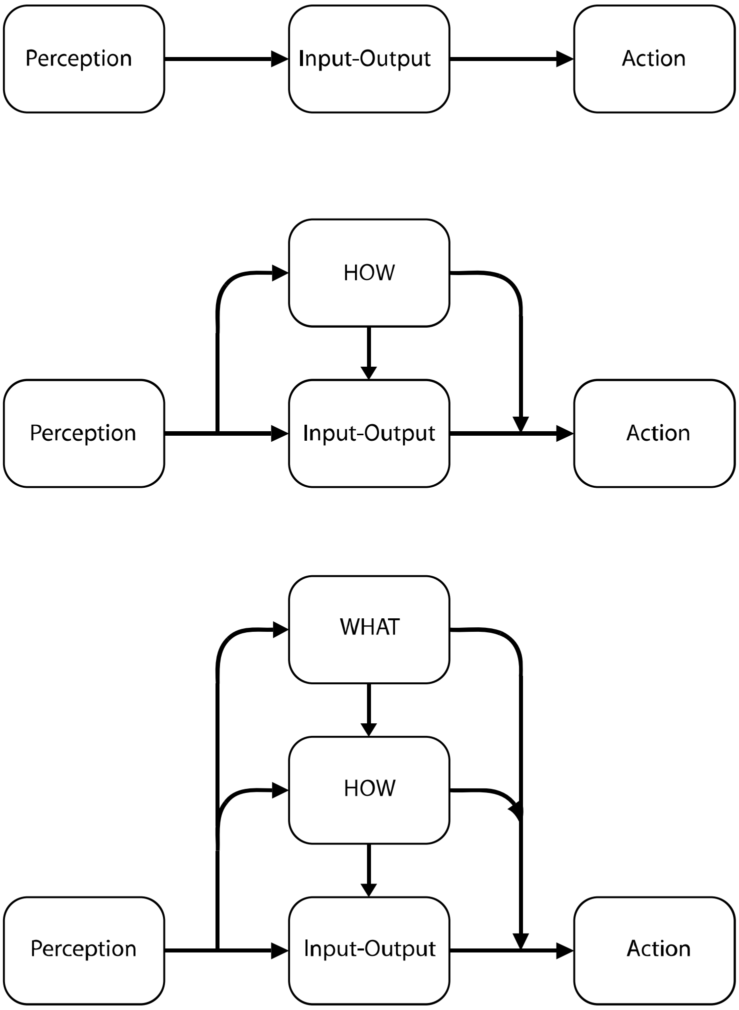

Figure 1. Three possible architectures: top) both *what* and *how* the system does is defined *a priori*; middle) the system modifies *how* it behaves but not *what* is doing; bottom) the system modifies both *what* and *how* it does.

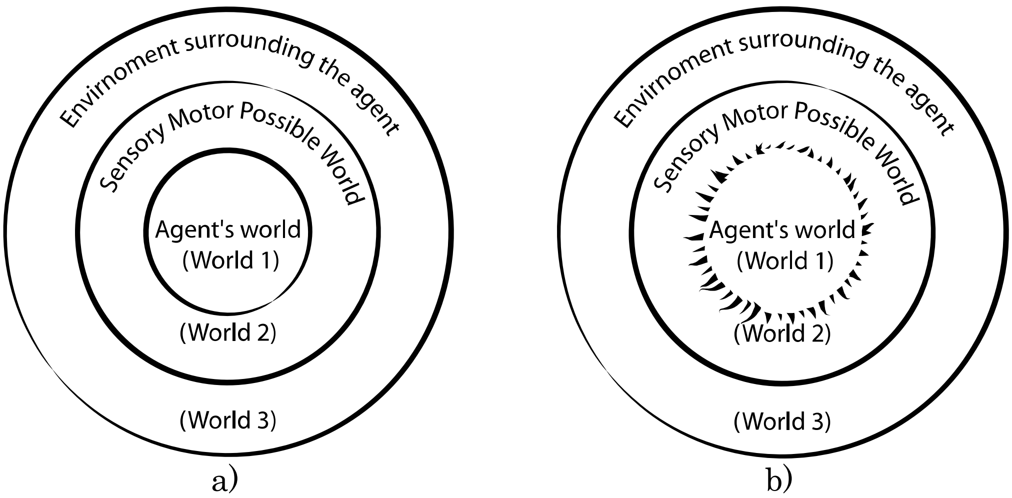

Figure 2. a) The agent's world is a subset of events of all the possible patterns and combinations of sensory motor stimuli that are a product of the interactions with the environment. In turn the environment is an even larger set of events. b) The agent's world can expand.

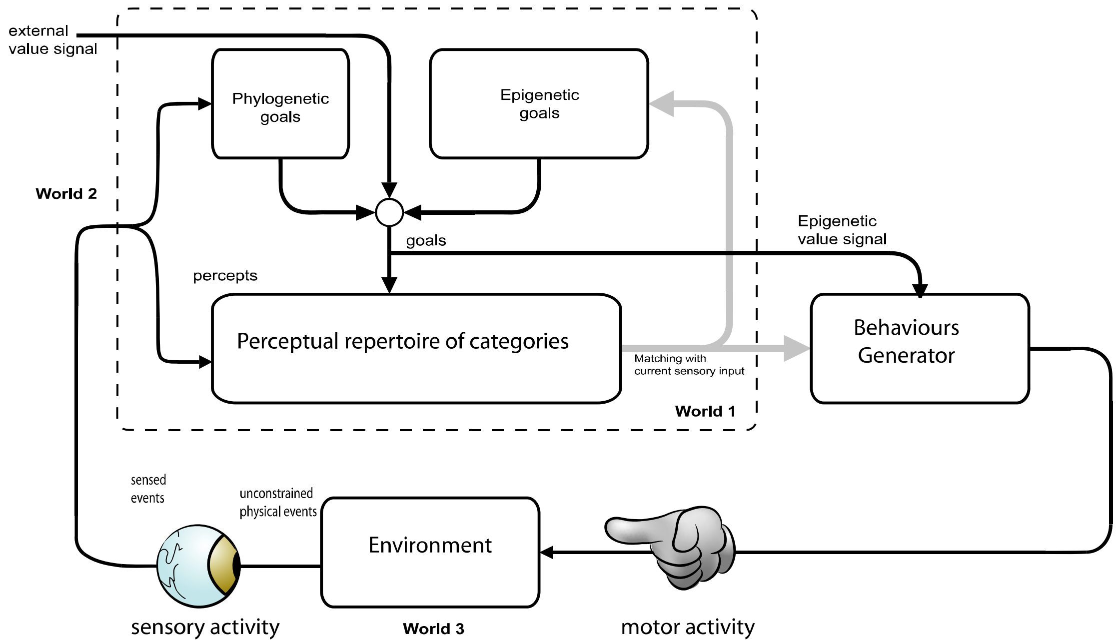

Figure 3. A general architecture capable of goal generations. Goals are dependent on the perceptual repertoire which is in turn responsible for the events among which goals are selected. It is the kind of circularity which is responsible for the self emergence of goals.