Multilevel anticipative interactions for goal oriented behaviors

Jean-Charles Quinton[1] quinton@n7.fr

Jean-Christophe Buisson[1] buisson@n7.fr

[1]University of Toulouse (INP/ENSEEIHT)
Computer Science Research Institute (IRIT)
2 rue Camichel 31071 Toulouse, France

Abstract

This paper suggests a possible infrastructure that can be used to build a cognitive architecture, combining several principles that appeared through evolution. Anticipation, regulation and coordination of interactive processes are used to produce goal-oriented behaviors, illustrated by a navigation application. Implications of the model for the phylogenesis to ontogenesis transition are discussed and a learning algorithm is presented as a perspective.

1. Introduction

Our work is mainly inspired by the Interactivist framework (Bickhard, 1993, Bickhard, 1996), which itself takes its roots in Piaget's constructivism (Piaget, 1952). It promotes a naturalistic and evolutionary approach which requires a drastic shift from a particle to a process metaphysics. We are also very interested in the Enaction paradigm (Varela et al., 1993), especially in the notion of life as based on conservation and adaptation of the organization through autopoiesis.

The main principles appearing during species evolution that we retain in this paper are self-maintenance through continuous regulation, synchronization with the environment, anticipation of the consequences of actions and coordination of processes (Quinton et al., 2008). Activity is necessary and central, in that only action and anticipation provide a normative knowledge with a direct epistemic contact with the environment (Bickhard and Christensen, 2002).

In the following model, knowledge, confidence or beliefs are all based on activity, as a synthesis of internal and external dynamics. Our approach can therefore be assimilated to situated and distributed dynamical systems or complex systems. Jun Tani detailed the theoretical and practical advantages of such approaches using recurrent neural networks (Tani, 2003). Similarly, the dynamic systems approach advocated by Esther Thelen from a psychological point of view reflects the same preoccupations (Thelen and Smith, 1994).

2. Implemented model

The computer-implemented model described in the following sections takes advantage of the previously cited principles to simulate a cognitive agent. Whether the agent is a robot acting in the physical world or a program communicating with a virtual environment, it interacts through a perception/action interface. All inputs and outputs must take real values but might be simple commands for rotating a given joint as well as complex features extracted from the optical flow.

2.1 Space

To integrate all information in one single homogenous system, we use a generic vector space whose dimensions correspond to the various perceptions and actions (Quinton and Inamura, 2007). To model higher functions of cognition and reflexivity, additional dimensions may be internally added during the agent's lifetime. In our model, all dimensions are derivative from the innate structure and interface of the agent's body. Still, the constructed dimensions may be associated with abstract processes, weakly coupled with the environment.

2.2 Interaction state

The basic elements or building blocks of the cognitive agent are single point states in any subspace of the previously described vector space. Additionally, an activity level is associated to each of these points. The more dimensions are set, the more specific is a state, corresponding to a particular or even unique event. Comparing states involves computing on their defined dimensions, thus defining a local environment based on the particular subspace. Com-

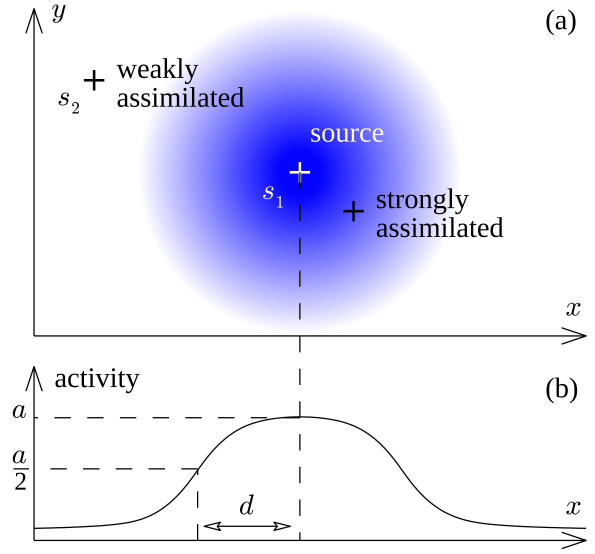

Figure 1: A 2D state defined by its x and y dimensions propagates its activity. Though the propagation is theoretically infinite, activity rapidly drops down, depending on the mean propagation distance d. Whatever are d or the source activity a, another state close to the source will be better assimilated than a further point, assuming Piaget's terminology.

mitting with a process metaphysics, states are not inert particles but interact with their local environment through an activity field. This field is represented on figure 1(a) for a two dimensional state. The equivalent effect from state s_1 to s_2 along dimension x is also plotted (b) and is defined by the propagation function in equation 2 where a is the source activity from s_1. The similarity function used is defined by the following sigmoid where the exponential argument coefficient of 5 is functional but arbitrary:

$$sim(s_1, s_2) = 1 - \frac{1}{1 + e^{-5\left(\frac{|s_2 - s_1|}{d} - 1\right)}} \quad (1)$$

$$p(s_1, s_2) = a * sim(s_1, s_2) \quad (2)$$

Extending this equation to multidimensional states s_i in subspaces S_i, the sigmoid is applied on the intersection subspace $\bigcap_i S_i$. As mentioned in the previous section, the interaction activity level, abstracting its input dimensions into one real value, can then be used as another dimension, undistinguishable from the original ones.

Our non classical approach is to merge perception and action in a single interactive state, then combining such points in an emergent dynamical system, interpolating actions as well as perceptions. The states and propagated activities continuously shape a global dynamical landscape which defines the agent's behavior.

2.3 Anticipation

What is called an anticipation is a special kind of interaction which gets part of its inputs from a distant state point, from which a simple diffusion would not bring any activity given the standard d parameter (see figure 1). To put it differently, anticipations are like instantaneous shortcuts in the activity landscape, spreading activity over large distances. The source and target states are in general in different subspaces. Using common usage terminology, an equivalent formulation for the most useful anticipations subclass would be: given the perceptive situation p_1, the actions a_1 will (certainly) lead to predicted situation p_2. The exact same anticipation can be defined as $[p_1, a_1] \rightarrow [p_2]$ or $[s_1] \rightarrow [s_2]$.

These pieces of knowledge can therefore be either passive as for an observer looking at a falling ball ("I'm here and will be there in the interaction space") or active as when pushing on something and anticipating the object touch feedback ("I'm doing that here and will go there"). Combining activities propagated to their source and target interactions (equation 3), these anticipations disturb the homogeneity of a state space else exclusively composed of interactions. The computations involved to integrate all the propagated activities represented on figure 2 at both ends of the anticipation arrow are detailed in equations 4 and 5.

$$a = (1 - \alpha) \times a_{src} + \alpha \times a_{tgt} \quad (3)$$

where $\alpha \in [0; 1]$ determines the trade-off between reactivity and anticipation in the system.

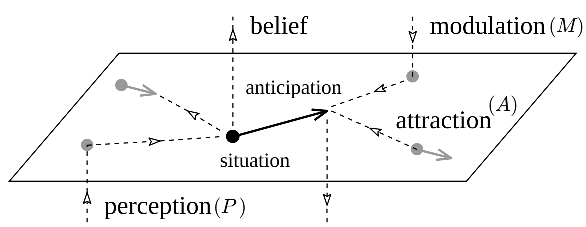

Figure 2: An anticipation integrates and propagates various activities. The figure focuses on one anticipation (black arrow), though all activities are updated in parallel. Dashed arrows represent activity propagation to and from the central anticipation. For instance, all anticipations (grey arrows) back-propagate their activity depending on the distance between their initial situation and target situation of other anticipations (attraction).

$$a_{src} = \max_{s \in M \bigcup A} p(s, src) \tag{4}$$

$$a_{tgt} = \max_{s \in P} p(s, tgt) \tag{5}$$

with M the set of top-down modulating activities, A the set of back-propagated activities from others anticipation sources and P the set of bottom-up perceptual activities.

These future oriented anticipations, by defining the direction of time, also break a deadly symmetry for the system, otherwise ineluctably reaching equilibrium states, no more differentiating perceived situation and anticipations. Without such elements, we need to split propagated and back-propagated activities, losing many immediate properties. Then we would also need to solve the classical dualistic problem of top-down and bottom-up integration at a global scale, where local activities are meaningless.

Given the constraints on the merging, integration and propagation functions provided below (of course verified by the previously introduced functions), activity can be trivially demonstrated to never diverge outside its initial range, anywhere in the interaction space (by combining the equations to determine an upper bound for the positive activity).

$$\forall (s_1, s_2) \in \mathcal{I}^2, p(s_1, s_2) \leq a_{s_1} \tag{6}$$

$$\forall a \in \mathcal{A}, a_{src_a} \leq \max_{s \in \mathcal{I}} p(s, src_a) \tag{7}$$

$$\forall a \in \mathcal{A}, a_{tgt_a} \leq \max_{s \in \mathcal{I}} p(s, tgt_a) \tag{8}$$

$$\forall a \in \mathcal{A}, a_a \leq \max\ (a_{src}, a_{tgt}) \tag{9}$$

where $\mathcal{A}$ is the set of anticipations and $\mathcal{I}$ the set of propagated interaction states.

2.4 Layer

This additional concept is not mandatory for the system to work correctly but useful to limit the complexity of interactions. Although we will use layers of interactions in the rest of the paper, the reader may keep in mind that a single layer with distant highly connected clusters is almost equivalent. Indeed, layers can be modeled by an additional dimension with specific anticipations replacing modulations between layers.

The further the layers are from the environment in terms of the anticipation path linking interactions, the more abstract and stable they are. Since they are loosely coupled with the environment, they do not depend on its fast and chaotic fluctuations (as in the optical flow), but rather rely on slow paced internal activity integrating lower level features.

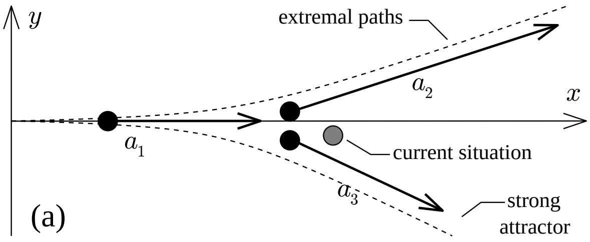

anticipations $\rightarrow$ a_1 a_2 a_3

$$\begin{bmatrix} c_1 \\ c_2 \\ c_1 \\ c_2 \end{bmatrix} * \left(a_{a_1} \times \begin{bmatrix} 1 \\ \\ 4 \\ \\ \end{bmatrix} + a_{a_2} \times \begin{bmatrix} 5 \\ 1 \\ 9 \\ 4 \end{bmatrix} + a_{a_3} \times \begin{bmatrix} 5 \\ -1 \\ 7 \\ -3 \end{bmatrix} \right) = \begin{array}{c|c} \begin{bmatrix} 4.8 \\ -0.8 \\ 7.1 \\ -2.2 \end{bmatrix} & \begin{matrix} x_s \\ y_s \\ x_t \\ y_t \end{matrix} \end{array}$$

with $a_{a_1} = .06$, $a_{a_2} = .13$, $a_{a_3} = .97$, $c_1 = \frac{1}{a_{a_1} + a_{a_2} + a_{a_3}}$, $c_2 = \frac{1}{a_{a_2} + a_{a_3}}$

(b) and $[m_i]^T * [n_i]^T = [m_i \times n_i]^T$

Figure 3: Collapse algorithm applied on three 2D anticipations. For each dimension, a barycenter is computed between anticipations with this dimension defined. Though an infinite number of paths is possible between anticipations a_2 and a_3, the resulting anticipation is strongly biased due to an attractor activity back-propagated towards a_3. As time passes, a_3 will increase its own assimilation of the situation, its activity and thus its influence on the taken actions.

2.5 Promoted actions

As stated before, perceptions and actions are entangled, and the resulting activity propagated by anticipations produces a field of activity shaping a behavioral landscape. Still when interacting with the real world, this superposition of potentialities has to be instantiated by unique and precise actions. To achieve it, all anticipations from a layer are collapsed into a single one by using a weighted sum pondered by activity (figure 3). As long as the input activities are continuous, the interpolation performed guarantees a smooth trajectory.

The resulting anticipation provides both the believed current situation and a target interaction state used to modulate lower layers. This target state keeps the agent moving towards attractors and future interactions by definition. The modulation interpolated at the lowest layers leaves the internal milieu of the agent's mind to be applied on the external dynamics. The system therefore dissolves the so called symbol grounding problem by propagating activities which only have meaning relatively to the local environment of interactions.

3. Navigation application

This application introduces a metaphor between physical paths and networks of anticipations that illustrates some underlying properties of the model. First of all, the simple fact that an agent can go by plane to London is not interesting if not her final destination. Suppose she lives in France and wants to go to Japan, going through England is not the shortest way. Yet it might be impossible to go directly to Japan from her place and having a connection in London might be the usual and even fastest way to go there.

The agent then has to integrate different pieces of knowledge to connect her flights. It is like reading personalized traffic signs with heuristic weights recursively computed depending on your final destination. This exact behavior is an emergent feature of the described model, propagating a higher activity in areas leading to goals through a sequence of known anticipations.

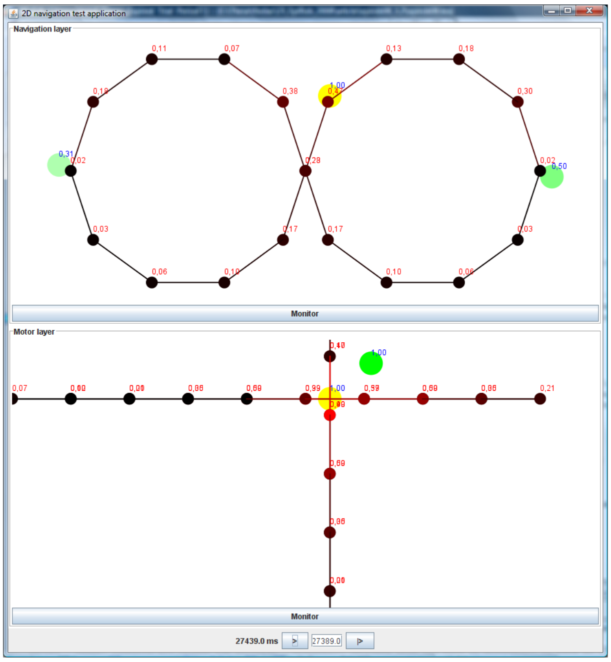

Figure 4: Graphical interface of the navigation program. The interaction space only possesses a minimal number of dimensions and a small set of anticipations for interesting behaviors to emerge. Though there might be more than screen coordinates dimensions, the interface only represents a 2D projection of the higher dimensional space. The navigation layer (left panel) modulates the motor layer (right panel) that interacts directly with the simulated environment. Perceptions are represented as yellow dots whereas modulations appear as green dots.

3.1 1-layer version

This first version is only introduced to show the limitations resulting from a single layer regulatory system. Anticipations can be represented as $[x_1, y_1, \Delta x_1, \Delta y_1] \rightarrow [x_2, y_2, \Delta x_2, \Delta y_2]$. Cartesian coordinates serve for both perceptions (x,y) and actions (Δx,Δy). Though these are not at all realistic dimensions for biological systems, they may still be grounded in robotic devices. A brief introduction to a more realistic perception using a similar model is given in (Basille et al., 2007).

At launch, though no special algorithm is implemented, an initialization phase of the interaction space takes place. Random initial activities rapidly converge to form a stable landscape synchronized with the environment. Only a few update cycles are necessary to spread activities through the entire network since propagation is instantaneous.

The internal and external dynamics then continuously interact through interpolated actions and perceptions, roughly following the oriented loop trajectories formed by the anticipations (as shown on figure 4 upper panel). This kind of large scale and long term anticipations are indeed only roughly satisfied since an interpolation of actions at this level is not powerful enough to provide a fine regulation on each effector.

3.2 2-layer version

Though the previous version exhibits a rough approximation of the targeted behavior, an additional layer is needed to improve the results. Two layers are therefore introduced and interconnected through modulation and perception. A navigation layer comprises positional anticipations of possible moves like $[x_1, y_1] \rightarrow [x_2, y_2]$ and a motor layer is composed of anticipations in direct relation with the environment such as $[x_1, \Delta x_1] \rightarrow [x_2, \Delta x_2]$ or $[y_1, \Delta y_1] \rightarrow [y_2, \Delta y_2]$.

The reader akin to constructivist theories may notice that the motor layer anticipations are more general than those from the navigation layer in that they are likely to be applicable on a wider range of situations and satisfied more often. Yet at the same time they are more specific in the sense that they deal with simple movements and a single dimension (adopting a more neuroscientific definition of specialization).

The inter layer modulation supplies the motor layer with a reachable target position as long as the navigation network is synchronized with and adapted to the environment. It bridges the gap between the different domains and scales on which the layers operate. Moreover, even if the x and y dimensions are directly mapped between layers here, the navigation layer may modulate any number of arbitrary joint specific layers. Thus the system can exhibit flexible and complex regulations as well as provide a first step towards generalization and body abstraction.

3.3 Dynamical goals

The main advantage of the presented approach is its implicit and continuous coordination of the anticipations. The most significant outcome is the possibility to have dynamical goals without having anything else to do than just changing or adding modulations at runtime. The modulation dynamics may be determined by another layer, the human user or even the environment with some kind of "virtual hormone".

The interaction space of figure 4 is coupled with its environment by standard perceptions and actions plus two modulations signals (displayed on the sides of the upper panel). The produced trajectories can be assimilated to those of the Lorentz's strange attractor, the missing dimension being the perception the modulation signals affect. The system keeps looping on one side until the fragile equilibrium between the two modulations has been reversed, the transition always taking place near the central implicit connection.

The way we articulate the simultaneously incompatible chewing, swallowing and breathing behaviors is a good biological example of such cycles and alternations. Though these actions seem trivial to us, they involve complex regulations depending on the texture of food, earlier physical efforts or even potential throat pain preventing the eater from swallowing as always.

4. Discussion

4.1 Comparison with existing approaches

The approaches referenced below have been selected as they share a lot of similarities with the model developed in this paper. Yet none of them in our view is equivalent nor implement full implicit coordination and regulation principles.

For instance, artificial recurrent neural networks for navigation and movement control (Tani, 1996) are also based on spreading activity but lack the structure flexibility resulting from local diffusion. The same difference exists with the hierarchical temporal memory model (Hawkins and Blakeslee, 2005). However, it not only uses activity as a unified form of representation, but also computes activities based on anticipation sequences. Similarly, interconnections between modules in the subsumption architecture need to be finely tuned (Brooks, 1999). Several other approaches are based on the coordination of processes but rely on explicit communication. Such approaches include the Polyscheme architecture (Cassimatis, 2005) or cooperative multi-agent system theories (Camps et al., 1998).

Compared to internal model based architectures (Haruno et al., 2003, Kawato, 1999), our behaviors are regulated at the emergent network level rather than relying on inner corrections of the forward and backward models. Interactions have no direct regulatory effect outside the field of their own perceptions and actions. Finally, partially observable Markov decision processes and equivalent statistical models are by essence sequential rather than parallel, taking a totally different perspective on phenomena.

4.2 Properties of the model

Since fundamental properties of the model are highly resistant to parameter variations and equation changes, basic elements can be easily extended with additional features. Integration over time for example improves robustness by simply adding an inertial factor in the merging equation. Though it might seem surprising or contradictory to the reader, the resulting oscillations represented on figure 5 also improve stability.

Time is not directly measurable as a mind "input" contrarily to many sensations. Moreover the difference has to be made between physical time (for example needed for neural propagations) and psychological time as perceived by the subject (Dennett, 1996). In our model, timing comes from the coupling with the environment, thus deriving from the laws of physics. Although adding an ad hoc objective perception of time is unrealistic, a solution coherent with the current model would be to synthesize central pattern generators by using bidirectionally coupled inertial anticipations. Combined with the current model, they could simulate rhythmic movements with an appropriate timing (Kuo, 2002).

In addition to the implicit time representation, several other aspects in our model are supported

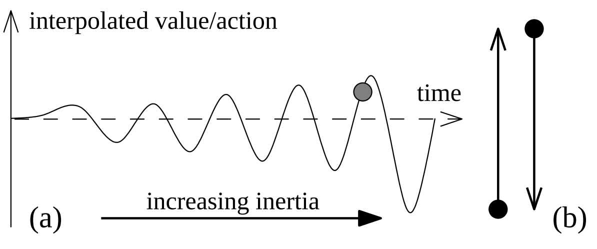

Figure 5: Behavior of an inertial anticipation. (a) When the inertia factor β added to equation 3 ($a_t = (1 - \alpha - \beta) \times a_{src} + \alpha \times a_{tgt} + \beta \times a_{t-\Delta t}$) increases in a low range, the system constantly probes the surroundings of any equilibrium point and oscillates around it, generating informative activity variations for the network. (b) Representation of the two concurrent anticipations interacting and promoting the action/position plotted on the left curve.

by experimental data in neuroscience. More specifically, the model might provide a plausible approximation to neuromodulation (Dehaene et al., 2006), brain wave synchrony (Engel et al., 2001), motivation concepts (Berridge, 2004), cortical layers hierarchy (Hawkins and Blakeslee, 2005) or mirror neurons (Rizzolatti et al., 2001).

Delimiting highly connected specific sub-networks, layers are an attempt to find a golden mean between pure modular integrative systems and giant homogenous networks. A full body motion is indeed required to compensate a single limb movement and keep balance, but this global regulation can be distributed between specific local networks. Though all interactions are more or less indirectly connected, they only strongly impact on a subset of dimensions.

More than just embedding regulations, the core of the model also accounts for goal-oriented behaviors. Interactions partially assimilating the environment or modulations but unable to satisfy their anticipations will continuously influence their surrounding network to be fully satisfied. With such a definition of goals and converging with recent researches in the domain (Oudeyer et al., 2007, Schmidhuber, 2006), an agent will develop curiosity for non totally assimilated situations but no interest at all in non assimilable situations.

Redundancy, rather than being a nuisance to be avoided or eliminated, improves the efficiency of the system. Since only the maximal propagated activity is integrated into an anticipation (equation 4 and 5), redundant anticipations have no effect on the activity landscape but will have more influence on the promoted actions if correctly assimilating the situation. This point is of great importance when the system cannot create any arbitrary compounds or when compounds cannot be aware of the global emergent function realized by the system (as for neurons relatively to the whole brain for instance).

In our model, there is an almost perfect symmetry between perceptions and actions, both are fundamentally entangled (partially accounting for mirror neurons interpretations). Yet there is also another symmetry between external and internal activities, both almost identically interacting with the network. Though the following reasoning may be adapted to hallucinations or dreams, we will describe a particular kind of optical illusions, namely perceptual filling-in. When a texture is partially erased but the rubbed out part is exactly projected on the blind spot of the retina by fixating a particular point, its seems like the human brain completes the pattern, whatever may be its complexity. Though the observer might perfectly acknowledge that his senses are deceiving him, the texture is perceived in its integrality. The brain does not really complete anything from our point of view, since perception is merely based on internal activity. Still, in the absence of visual sensations, anticipations only integrate back-propagated activity and modulations. Low level visual anticipations compatible with the surrounding pattern get an activity boost not overwhelmed by real sensations, therefore propagating the belief of a full texture to higher layers as consciously experienced.

Finally, though perception and action have symmetric roles, a distinction must be made between generation and recognition. Because activity and timing are provided by the environment during recognition, the internal dynamics can easily synchronize on any previously learned pattern. Though the structure and organization of the anticipations may not be different, generating the same pattern by only integrating modulation signals is unlikely. Only the most usual patterns that shaped the layer emerge from the spontaneous inner propagation of activity. Even if our model concentrates on low-level cognition, a similar gap exists between understanding and speaking abilities in the language domain.

4.3 From phylogenesis to ontogenesis

To focus the discussion on the specific theme of the conference, this section deals with the relations between behaviors inherited from phylogenesis and skills acquired during life. Firstly, we consider that learning appeared during evolution to allow organisms to adapt to their no longer genetically predictable environment. Indeed, an increasing number of possible interactions between an organism and its environment results in an infinitely more complex perceptual world.

Even if this aspect has not been emphasized in this paper until now, anticipations can be learned and confirmed by probing actions on the environment and verifying the satisfaction of the consequences. Even if such a mind was initially blank, the coupling with the physical body would rapidly generate a large number of passive anticipations. Reflexes and metabolism already provide a huge quantity of survival information selected through evolution, thus defining the laws governing the immediate environment of the agent's mind.

This basic set of anticipations could then progressively be extended by constantly interacting with the whole agent's environment. When building anticipations mainly based on internal activity, the coupling with the environment would get progressively weakened though they would still be grounded in perception and action. Depending on the development context of each individual, unique networks of activities would emerge. These would constitute stable attractors for the internal dynamics and influence all behaviors, defining habits, interests, obsessions, values and identity (Moreno, 2000). Trying to assimilate the situation at any time, they would apply if not in direct conflict with stronger attractors. For instance, behaviors totally useless for the survival such as taping a rhythm with the foot, chewing some non eatable thing or playing with a pen are common when our limbs are "free". Similarly, humming a sound is easy if not trying to talk or listen to other music.

This phenomenon might be correlated with the evolution of brain plasticity during development. In a baby's mind, innate behaviors are hardly connected and apply to different realms. Contradictory anticipations can coexist as not consciously perceived as such, since they are not linked with additional knowledge increasing the coherence of the system. Moreover almost everything is considered as new, so anything not too challenging can be learned and accepted. The behavioral dynamical landscape is shaped by the stimulations. Afterwards, the older the mind becomes the more structured and coherent is the network. An adult has more remembrances, masters more skills. On one hand, the interaction network is more stable and resistant to noise, thus the agent will be able to generalize situations correctly. On the other hand, it becomes harder to think of a situation as novel, since it is always vaguely assimilated to past experience. Thus accepting destabilizing changes or learning totally novel fields (without a lot of motivation, concentration and rehearsal) is threatened by the overall stability of the mind. Though change is always possible in such mathematically chaotic systems, a subtle combination of complementary successive influences is required.

Drives such as thirst, carefully selected by evolution for the sake of the whole body (mainly for the agent to be able to reproduce), are interacting with learned processes. Therefore they take part in the overall activity and indirectly influences cognitive actions such as grasping a bottle of water (that was not selected as good by evolution). In fact any interaction or behavior, as long as it is not hindering highly active processes will just coordinate with them. To put it differently, notions of goodness or badness, usefulness or futility are not absolute. The human drinking behavior has no immediate value for a single cell absorbing water molecules through its wall. As long as the environment provides the necessary elements and functional interactions are preserved, any self-maintenant system will just work properly. Breathing machines can for instance keep organs alive even after brain death.

We can draw a parallel between self-maintenant far from thermodynamic equilibrium systems, or autopoietic systems (Varela et al., 1974), and closed networks of interactions/anticipations. The promoted actions keeps the agent within the network as long as the environment remains in adequation. Although most of the discussion remains hypothetical and subject to experimental validation, it would be scientifically simple and economic to have the same principles range from the cell level to human cognition.

5. Conclusion and perspectives

The validity of the model now needs to be further tested on larger scale problems. Introducing a learning mechanism similar to the one presented at Epigenetic Robotics 2007 is required in order to produce genuine anticipations from a basic set of innate reflexes (Quinton and Inamura, 2007). A confidence reinforcement algorithm may select anticipations with highly correlated activity in the assimilation of the situation and in the satisfied target interaction. Such algorithm is however required to apply synchronously on several levels and adapt to various rhythms and scales. Moreover, due to the potential huge number of learned anticipations and taking intro account the sparsity of the interaction space vectors, the optimizations implemented last year need to be greatly improved.

Additional theoretical problems will certainly appear before reaching the flexibility of human level regulations as described by Merleau-Ponty (Merleau-Ponty, 1963). Adapting our model to control a simple haptic device so it can learn and interact in a real environment is however a first necessary step. Extensively testing our hypotheses will hopefully help us to detect the flaws in our theories and find the very principles underlying low-level cognition before going any further.

References

Basille, J.-L., Buisson, J.-C., and Quinton, J.-C. (2007). Interactivist navigation. In *Proceedings of ECSC'07 - European Cognitive Science Conference (Delphi, Greece)*, page 922. Lawrence Erlbaum Associates.

Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience. *Physiology & Behaviour*, 81.

Bickhard, M. H. (1993). Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence.*

Bickhard, M. H. (1996). The emergence of representation in autonomous embodied agents. In *AAAI Fall Symposium on Embodied Cognition and Action.*

Bickhard, M. H. and Christensen, W. D. (2002). Process dynamics of normative function. *Monist*, 85(1):3–28.

Brooks, R. A., (Ed.) (1999). *Cambrian Intelligence - The early history of new AI.* MIT Press.

Camps, V., Gleizes, M.-P., and Trouilhet, S. (1998). Properties analysis of a learning algorithm for adaptive systems. *International Journal of Computing Anticipatory Systems*, 1:223–233.

Cassimatis, N. L. (2005). Integrating cognitive models based on different computational methods. In *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*, pages 402–407.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., and Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*.

Dennett, D. (1996). *Kinds of Minds: Toward an Understanding of Consciousness*. The Science Masters Series. New York: Basic Books.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews*, 2.

Haruno, M., Wolpert, D. M., and Kawato, M. (2003). Hierarchical mosaic for movement generation. *International Congress Series*, 1250:575–590.

Hawkins, J. and Blakeslee, S., (Eds.) (2005). *On Intelligence*. Times Books.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*.

Kuo, A. D. (2002). The relative roles of feedforward and feedback in the control of rhythmic movements. *Motor Control*, 6:129–145.

Merleau-Ponty, M. (1963). *The Structure of Behavior*. Beacon Press, Boston.

Moreno, A. (2000). Closure, identity, and the emergence of formal causation. *Annals of the New York Academy of Sciences*, 901:112–121.

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*.

Piaget, J., (Ed.) (1952). *The Origins of Intelligence in Children*. International Universities Press.

Quinton, J.-C., Buisson, J.-C., and Perroto, F. (2008). Anticipative coordinated cognitive processes for interactivist and piagetian theories. In *Proceedings of AGI'08 - Artificial General Intelligence (Memphis, USA)*, pages 287–298. IOS Press.

Quinton, J.-C. and Inamura, T. (2007). Human-robot interaction based learning for task-independent dynamics prediction. In *Proceedings of EpiRob'07 - International Conference on Epigenetic Robotics (Piscataway, USA)*, pages 133–140. Lund University Cognitive Studies.

Rizzolatti, G., Fogassi, L., and Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2:661–670.

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

Tani, J. (1996). Model-based learning for mobile robot navigation from the dynamical systems perspective. *IEEE Trans. on Systems, Man, and Cybernetics Part B: Cybernetics*.

Tani, J. (2003). Symbols and dynamics in embodied cognition: Revisiting a robot experiment. *Anticipatory Behavior in Adaptive Learning Systems*, pages 167–178.

Thelen, E. and Smith, L. B., (Eds.) (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press.

Varela, F. J., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: the organization of living systems, its characterization and a model. *Biosystems*, 5:187—196.

Varela, F. J., Thompson, E., and Rosch, E., (Eds.) (1993). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.