# Interactivist sensorimotor learning: computational implementation and parallel optimization

Jean-Charles Quinton and Jean-Christophe Buisson

University of Toulouse, Computer Science Research Institute,
2 rue Charles Camichel, 31071 Toulouse, France
quinton@n7.fr,buisson@n7.fr

## 1 Background

We here introduce an approach where all knowledge is fundamentally grounded in interactions. We assume that perception and action are intertwined at any level and for any kind of behavior. Generation and recognition are therefore the two sides of a single process, with perception and action occurring simultaneously and continuously. This work is based on the interactivist framework developed by Bickhard et al. [1, 2], taking a naturalistic and evolutionary approach to cognition. In this sense, it shares similarity with enaction as developed by Varela et al. [3] and incorporate Piaget's constructivism [4].

We aim at understanding and implementing complex and generic regulation of actions. The model therefore focuses on the implicit modulation and coordination of processes. In our view, sensorimotor capabilities emerge from interactions and anticipations using an homogeneous form of representation, as a synthesis of internal and external dynamics. We are therefore akin to dynamic systems approaches as advocated by Esther Thelen [5] or Jun Tani [6]. Architectures similar to our own such as the subsumption architecture [7], the Polyscheme architecture [8] or the hierarchical temporal memory model [9] all rely on explicit communication or wiring to make the various elements cooperate (either called modules, codelets or experts). Internal forward and backward models [10] make use of inner corrections in the sensorimotor loop to regulate their actions, but the global system has to select a single module that will actually perform in the environment, therefore introducing an unbridgeable gap between the various potential behaviors. Artificial neural networks and other parallel distributed processes approaches based on spreading activity often draw a clear cut between diffusion and point to point exchanges. Our proposal is an attempt to merge the two forms of communication to achieve online and implicit coordination.

## 2 Principles

The whole theory and applications gravitates around the central notion of interaction. All other principles developed in the following paragraphs (i.e. regulation, coordination, learning) derive from it. Interactions can be found as soon as coupled dynamical systems are communicating, whatever their complexity. If we consider one system as the subject and the rest as its local environment, acting is like asking questions by influencing the environment dynamics and perceiving consists in listening to the answer. This feedback may be as simple as a touch sensation when probing surrounding objects or as complex as an uttered sentence from a human being.

To keep being adapted to the situation, interactions need to be flexible enough to cope with noise and variability. Since interactions have a narrow view on the situation (defined by their perceptions and actions), their instant adaptation is never absolute. Rather than relying on a centralized control of discrete behaviors, a second level of interactions may occur between the various specialized processes. The agent is therefore an emergent system resulting from the spread out activity of its component processes.

In a real world where physical laws rule the global dynamical landscape, being adapted implies being synchronized in time and space with the surrounding environment (metabolism is for instance synchronized with the day/night alternation). For internal processes not directly coupled with the external world, this is somehow similar to coordinating with other processes within what is often called the internal milieu. Such coordination leads to a dynamical goal/means subordination, where a goal acts as an attractor defined by active interactions not perfectly adapted to the situation. Indeed, interactions being subject to delays in sensorimotor loops, consequences of actions need to be anticipated; though the situation might be partially assimilated, corresponding actions may not be possible and consequences not immediately satisfied (grasping a visible door knob that is still out of reach for example).

Learning is necessary as soon as the agent's environment is no more phylogenetically predictable. Since the agent enacts its own world through its senses, this aspect is more correlated to motor and perceptive capabilities than to the environment complexity. Although the history of the species determines many physical and metabolic characteristics, a huge number of behavioral degrees of freedom remains. Only agents with learning capabilities can therefore shape their future dynamics by relying on past individual experiences.

Since interactions occur at all times and impact on both the internal and external dynamics relatively to the agent, online learning is required to continuously adapt to the ever changing dynamical coupling. In a system where the anticipation of consequences of actions is crucial, we promote an unsupervised variation/selection learning mechanism guided by the relative assimilation level of interactions. Parts of the interaction network that cannot be maintained active are progressively discarded, while new interactions continuously try to assimilate the dynamics. From a local perspective, sub-networks survive or die by strengthening or losing the bonds between their elements, thus also locally using a reinforcement learning approach.

## 3 Computational model

Interactions are represented in a mathematical vector space where dimensions stand for any perceptual stimulus, motor command or internal signal. This unified representation is combined with an activity field to represent the overall dynamics. Activity is propagated through the field using variants of inverse distance weight functions as introduced by Shepard [11]. Such a homogeneous field could not account for any learning or even reactive behavior without additional elements. Thus anticipations are introduced as activity shortcuts between distant points (figure 1a). They integrate top-down modulations and bottom-up sensations to define the possible sensorimotor behaviors. At the emergent network level, the actions actually performed by the agent are continuously interpolated as a weighted sum in which highly adapted interactions correspond to the strongest attractors. Although this is only a short description of our previously implemented model [12], we detail here its strength in terms of scalability and an ongoing extension developed in order to account for learning.

Depending on the coupling history and the current dynamics, each anticipation can reach three extreme situations:

- Both the context (including taken actions) and predicted consequence may be correctly assimilated and the corresponding interaction needs to be reinforced.
- The context did not match the situation and the anticipation was just not adapted, like an old remembrance that was once true.
- Only the consequence was not satisfied. In this case, either hidden variables need to be invoked as described by Perotto et al. [13] or the anticipation is just incorrect and needs to accommodate.

A simple mechanism to account for all intermediate situations is to introduce a lifespan for each anticipation. This would be a generalization of the confidence level we already used for sensorimotor exploration [14]. Correlated activity in the context and consequence of anticipations would lead to a greater lifespan and reciprocally unadapted anticipations would be progressively eliminated, introducing at the same time a forgetting ability within the system. Moreover redundant anticipations are not a hindrance to be avoided or controlled but rather reinforce the adapted activity and confirm the predictions, thus allowing the introduction of new (and even random) anticipations at runtime. To make a long story short, only coordinated anticipations participating in an emergent active network would survive (figure 1b). Though a single anticipation would not define any property or concept and could be easily replaced, it would still be a building block for more complex behaviors. This is to some extend similar to Varela's discussion about the tessellation automaton in autopoiesis theory [15], where cell components are regularly produced to regenerate the membrane.

Though we did not solve large-scale problems until now, we refined the model in toy applications. These were used to test our hypotheses on the roots of cognition and determine the basic set of principles required to perform general sensorimotor tasks in real environments. Learning algorithms were introduced in appli-

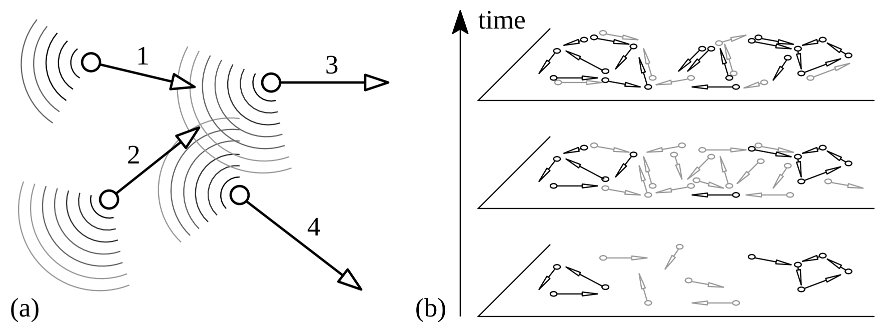


**Fig. 1.** *(a)* Anticipations (*arrows*) propagate activity through the interactive space. Only attraction toward goals is displayed (*gray arcs*), *3* is acting as an attractor for *1* and *2*. *(b)* New interactions (*gray arrows*) are reinforced by the coupling with the environment. Stable networks are produced and redundant anticipations generated.

cations and simulations ranging from rhythm tracking [16] to mechanical device manipulation [14]. Whether the programs were using pure mutation/selection or reinforcement mechanisms, human users could interact without any calibration or noise reduction in the input signals.

Even if most of the tasks introduced by the applications are easily and naturally performed by humans, such as reaching an object, they often involve a complex dynamics and many degrees of freedom. To allow a progressive construction of networks of anticipations and fasten the computations, sparse vectors have been used to represent the interactions. Additionally and contrasting with many existing models of sensorimotor behaviors, this approach is easily scalable and adaptable to any heterogeneous substrate. As long as the interactions are preserved, allowing the bidirectional propagation of activity between distant modules, the interpolation necessary to decide for actions can be hierarchically computed.

Moreover, for such an approach to be efficient on larger scale problems, parallel implementations and optimizations of our model may be required. Multi-core architectures are becoming standard nowadays and this tendency is reflected by the amount of work put in the development of generic abstraction libraries such as OpenVidia [17] or Brook [18]. Though software often needs to be completely rethought to be translated into parallel code, the explicit scheme of the finite difference method used for updating the internal dynamics as well as the interpolation algorithms used in our model are particularly well suited for distributed processing. Taking advantage of the power of graphics processing units available on any modern personal computer, a parallel implementation of the model has been developed using shader programs. Additionally and to adapt to the new generation of distributed architectures, a Cell Broadband Engine version has been studied and the model was partly ported on the Playstation®3 Linux operating system.

## References


1. Bickhard, M.H.: The emergence of representation in autonomous embodied agents. In: AAAI Fall Symposium on Embodied Cognition and Action. (1996)
2. Bickhard, M.H., Christensen, W.D.: Process dynamics of normative function. Monist **85(1)** (2002) 3–28
3. Varela, F.J., Thompson, E., Rosch, E., eds.: The Embodied Mind: Cognitive Science and Human Experience. MIT Press (1993)
4. Piaget, J., ed.: The Origins of Intelligence in Children. International Universities Press (1952)
5. Thelen, E., Smith, L.B., eds.: A Dynamic Systems Approach to the Development of Cognition and Action. MIT Press (1994)
6. Tani, J.: Model-based learning for mobile robot navigation from the dynamical systems perspective. IEEE Trans. on Systems, Man, and Cybernetics Part B: Cybernetics (1996)
7. Brooks, R.A., ed.: Cambrian Intelligence - The early history of new AI. MIT Press (1999)
8. Cassimatis, N.L.: Integrating cognitive models based on different computational methods. In: Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society. (2005) 402–407
9. Hawkins, J., Blakeslee, S., eds.: On Intelligence. Times Books (2005)
10. Haruno, M., Wolpert, D.M., Kawato, M.: Hierarchical mosaic for movement generation. International Congress Series **1250** (2003) 575–590
11. Shepard, D.: A two-dimensional interpolation function for irregularly-spaced data. Proceedings of the ACM National Conference (1968)
12. Quinton, J.C., Buisson, J.C.: Multilevel anticipative interactions for goal oriented behaviors. In: Proceedings of EpiRob'08 - International Conference on Epigenetic Robotics (Brighton, UK), Lund University Cognitive Studies (2008)
13. Perotto, F.S., Buisson, J.C., Alvares, L.O.: Constructivist anticipatory learning mechanism (calm) dealing with partially deterministic and partially observable environments. In: Proceedings of EpiRob'07 - International Conference on Epigenetic Robotics (Piscataway, USA), Lund University Cognitive Studies (2007) 117–124
14. Quinton, J.C., Inamura, T.: Human-robot interaction based learning for task-independent dynamics prediction. In: Proceedings of EpiRob'07 - International Conference on Epigenetic Robotics (Piscataway, USA), Lund University Cognitive Studies (2007) 133–140
15. Varela, F.J., Maturana, H.R., Uribe, R.: Autopoiesis: the organization of living systems, its characterization and a model. Biosystems **5** (1974) 187–196
16. Buisson, J.C.: A rhythm recognition computer program to advocate interactivist perception. Cognitive Science **28**(1) (2004) 75–87
17. Fung, J., Mann, S.: OpenVIDIA: parallel GPU computer vision. In: MULTIMEDIA '05: Proceedings of the 13th annual ACM international conference on Multimedia, New York, NY, USA, ACM (2005) 849–852
18. Buck, I., Foley, T., Horn, D., Sugerman, J., Fatahalian, K., Houston, M., Hanrahan, P.: Brook for GPUs: stream computing on graphics hardware. In: SIGGRAPH '04: ACM SIGGRAPH 2004 Papers, New York, NY, USA, ACM Press (2004) 777–786