

<http://his.diva-portal.org>

This is an author produced version of a paper presented at The fourth workshop on Anticipatory Behavior in Adaptive Learning Systems, ABiALS, 26 June 2008 in Munich. This paper has been peer-reviewed but may not include the final publisher proof-corrections or pagination.

Citation for the published paper:

Montebelli, Alberto; Lowe, Robert & Ziemke, Tom

"The Cognitive Body: From Dynamic Modulation to Anticipation"

Anticipatory Behavior in Adaptive Learning Systems, 2009 (Lecture Notes in Computer Science vol 5499) pp 132-151

ISBN: 978-3-642-02564-8

URL: [http://dx.doi.org/10.1007/978-3-642-02565-5\\_8](http://dx.doi.org/10.1007/978-3-642-02565-5_8)

Access to the published version may require subscription.

Published with permission from: Springer Verlag

# The Cognitive Body: from Dynamic Modulation to Anticipation

Alberto Montebelli, Robert Lowe, Tom Ziemke

University of Skövde  
School of Humanities and Informatics  
SE-541 28 Skövde, Sweden

{alberto.montebelli, robert.lowe, tom.ziemke}@his.se

**Abstract.** Starting from the situated and embodied perspective on the study of cognition as a source of inspiration, this paper programmatically outlines a path towards an experimental exploration of the role of the body in a minimal anticipatory cognitive architecture. Cognition is here conceived and synthetically analyzed as a broadly extended and distributed dynamic process emerging from the interplay between a body, a nervous system and their environment. Firstly, we show how a non-neural internal state, crucially characterized by slowly changing dynamics, can modulate the activity of a simple neurocontroller. The result, emergent from the use of a standard evolutionary robotic simulation, is a self-organized, dynamic action selection mechanism, effectively operating in a context dependent way. Secondly, we show how these characteristics can be exploited by a novel minimalist anticipatory cognitive architecture. Rather than a direct causal connection between the anticipation process and the selection of the appropriate behavior, it implements a model for dynamic anticipation that operates via bodily mediation (*bodily-anticipation hypothesis*). This allows the system to swiftly scale up to more complex tasks never experienced before, achieving flexible and robust behavior with minimal adaptive cost.

## 1 Introduction: a cognitive body

While there is much agreement that cognition is *embodied*, it remains less clear exactly what role(s) the body plays in cognitive processes. The obvious common-sense answer would highlight the role of the body in dictating the physical relation between an agent and its environment. Of course the fundamental function that the body of any organism plays - the appropriate adaptive situating of its available set of sensors and actuators in a spatio-temporal frame of reference - cannot be missed or ignored. However, this common-sense interpretation is broadened in embodied and situated approaches to the study of cognition, both at theoretical and experimental levels [1,2,3,4,5]. The body shapes the cognitive potential of the agent by completely specifying the nature and range of all possible interactions with its environment. A self-organized agent typically depends on and deeply exploits such constraints [6].

The basic idea of a highly systemic approach to the study of cognition was already centrally present in the work of early cyberneticists (i.e. [7,8]), gestalt psychology [9] and ecological psychology [10]. The sudden rise of cognitive science cast a shadow on such historically prominent intellectual work. This should not come as a surprise: apparently, large bodies of sometimes outstanding scientific knowledge are destined to be reconsidered, or even completely rediscovered, over and over, whenever there is an intellectual need for them. Presently massive research efforts investigate the problem of understanding cognition by a systematic decomposition. In the scientific tradition, reductionism proved powerfully effective in producing sound explanations (with predictive power) of natural phenomena. Nevertheless, it is wrong to infer that all explanations are reductionistic. Such a misconception might be particularly pernicious in an epoch where the scientific community masters, and has large availability of, the necessary technology to engage in the exploration of the problem of non-linear complexity, and is intellectually committed to the development of the appropriate mathematics to start addressing it. That is, dynamic systems theory offers a natural language to a systemic approach to the study of adaptive behavior and cognition [11]. Much work has already clarified the need for a consistent deployment of the existing mathematical tools and for their further development [12,2,13,14].

Nowadays, a more systemic view of the mind pervades at least a few major theoretical frameworks in the study of cognitive processes. Several authors are currently committed to the underpinning of a theoretical background, in which the specific embodiment of an organism has non-trivial cognitive consequences. The body massively pre-/post-processes the information flow to and from the nervous system, and the common evolutionary history and ontogenesis of body and nervous system provides a deep, distributed integration of bodily and nervous functions (e.g. see [15]). Perception and action are not causally sequential activities, but can be seen as closely interrelated and in fact inseparable, one supporting the other [16,6].

Nevertheless, we have reason to think that this perspective does not go far enough. Rather than treating the body as a mere interface to the world, we should also take into account what happens inside the body of an organism, and its potential cognitive consequences [17,18,19,20,21]. We find that the hidden, bio-regulatory dynamics developing under the surface of the body are largely neglected in the study of cognitive phenomena. As some authors put it, the interaction between bio-regulatory events that take place inside the body of an agent and what is traditionally interpreted as its control system, might be a crucial component of its ongoing cognitive processes [17]. In line with this thread, in this paper we describe our current experimental work in cognitive robotics, focusing on the role that the intrinsic non-neural bodily dynamics might play in supporting and boosting cognition. In Section 2, we discuss some preliminary results showing how a non-neural internal state can modulate the activity of a simple neurocontroller. We then formulate the programmatic foundations for an extension of our work towards a bodily-mediated anticipatory cognitive architecture. Firstly, on theoretical grounds, we advocate that non-neural bodily

dynamics might play a fundamental role in a new anticipatory cognitive architecture (Section 4). Secondly and more concretely, we report on the initial experimental analysis of this idea (Section 5). Then, we briefly comment on the theoretical implications of our work and on necessary future developments (Sections 6 and 7).

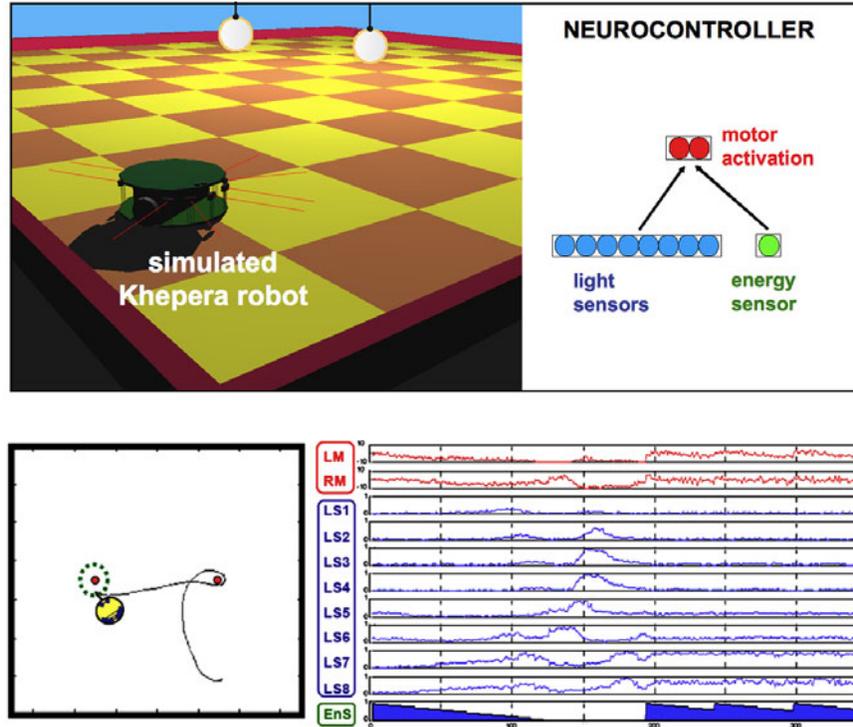
## 2 The dynamic role of the cognitive body: a minimalist case study

In a recent study [22,23], we have shown how even very simple non-neural bodily states can play a crucial role in the modulation of the activity of an artificial nervous system, i.e. on the behavior generated by an artificial neural network (ANN) implementing the neurocontroller of our simulated robotic agent. We used standard evolutionary algorithms to set the weights and biases for a simple reactive ANN with no hidden layers, driving the motoneurons of a simulated Khepera robot (see Figure 1). The system self-organized in order to find a recharger for its energy level (i.e. each instantiation of ANN during the evolution was simply rewarded for the maintenance of a positive level of energy, punished otherwise), thus overcoming its temporal linear energy decay. The invisible recharger was placed in a circular area centered under one of the two visually identical light sources, randomly selected for each replication. An energy level sensor, together with a battery of light and infra-red sensors constituted the sensory inputs to the ANN.

As part of the analysis of the successfully evolved system<sup>1</sup>, we manipulated the energy level as the control parameter for the whole system [12,24]. By systematically clamping<sup>2</sup> it to a discrete set of possible values, ranging from zero to 'full', we observed and classified a number of possible behaviors, exemplified in Figure 2- left. After exhaustion of the behavioral transients, we found three general classes of qualitative behavioral attractors. We observed: exploratory behaviors at the lowest levels of energy, i.e. the agent engaged in loops between potential energy sources and also in external loops broadening its explorations to the rest of the environment (i.e. see trajectory in panel 'A' of Figure 2); more local behaviors at higher levels of energy (i.e. the agent was closely looping in the neighborhood of a single source as in panel 'C'; hybrid behaviors, embedding characteristics from both previous classes (as in panel 'B') for intermediate levels

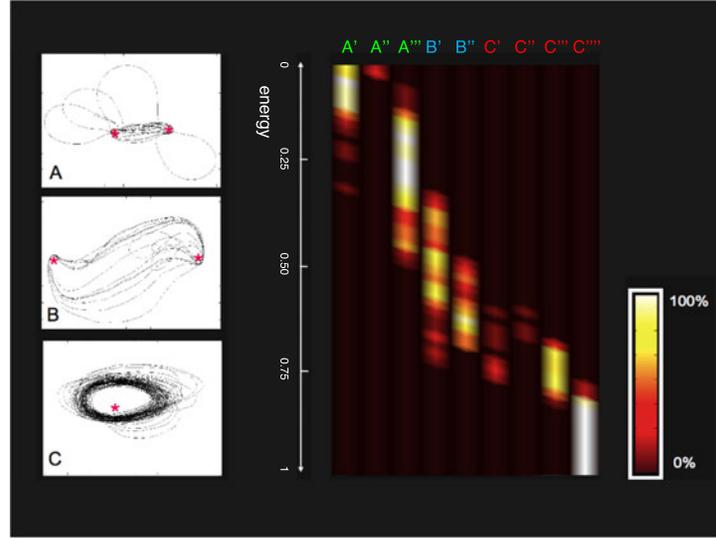
<sup>1</sup> By the term 'system', here and in what follows, we refer not just to the evolved ANN, but to the global system constituted also by the agent, its environment and its non-neural dynamic mechanism of the energy level. Therefore, cognition is here conceived and analyzed as a broadly distributed process; a cognitive aggregate, rather than a localized and proprietary process.

<sup>2</sup> The term 'clamping' here refers to the injection of a constant energy level as input for the ANN during the whole duration of the replication of the experiment. The agent is free to behave in its environment for a period of time sufficient to exhaust all behavioral transients and permit observations of satisfactory duration.



**Fig. 1. top** - Experimental setup. A simulated Khepera robot moves in a square environment containing two visually identical light sources (suspended white light bulbs). Its neurocontroller is a simple reactive ANN with no hidden layers, directly controlling the two motors of the robot and receiving information from the light sensors and from an energy level mechanism. The choice of such a simple scenario aimed to facilitate our analysis and emphasize the object of study. **bottom** - Example of the evolved behavior. As its energy level sensor measures the temporal linear decay (graph labeled EnS), the simulated agent (large cylinder) approaches the light to the right (filled circle). The recharging area (dashed circle) is invisible to the robot, that can sense it only by virtue of its effect on the energy level sensor. As its energy reservoir is instantaneously refilled to the maximum level, the agent is engaged in a stable behavior in the proximity of the rewarding light source. The signals labeled LM and RM show, as a function of time, the activation of the left and right motors; LS1-8 represent the activation of the light sensors.

of energy. The relative frequency of the three groups of behaviors was reliably dependent on the current energy level (Figure 2- right).



**Fig. 2.** **left-** Sample spatial trajectories for the three classes of behaviors observed in clamped conditions after transient exhaustion. Exploratory behaviors (panel A), local behaviors (panel C) and hybrid forms (panel B). Potential energy rechargers (i.e. the position of the light sources) are indicated by red stars. For a better resolution of details, the icons representing each class of trajectories zoom on the area of main interest surrounding the light sources. **right-** The intensity of the pixels for each column (corresponding to attractors belonging to classes A-C, as specified by their labels on the top row) represents the relative frequency of the behavioral attractor as a function of the energy level. For example, an energy level of 0.7 leads to the expression of attractor C''' (in 70% of the replications), C' (20%) or B' (10%). For energy levels in the interval  $[0.0, 0.4]$  we can observe a clear dominance of attractors in class A. A similar dominance in the energy interval  $[0.7, 1.0]$  is shown by attractors in class C. The hybrid forms in class B characterize intermediate energy levels. Adapted from [22]

Regarding the evolutionary task, we then examined the implications of the behaviors that we observed in clamped conditions. As the energy level is left free to follow its natural dynamics, it constitutes an effective self-organized and dynamic action selection mechanism. Different classes of behaviors are locally available to the agent as a function of its current energy level. Apparently, high energy levels imply that a source of energy was recently visited. Given the obvious physical constraints on the agent's speed, it follows that it must be still in the proximity of the agent, consistent with the selection of local behaviors. On the other hand, low energy levels imply that the recent search for an energy source was unsuccessful. This effectively correlates with broader exploratory be-

haviors. The solution of this minimalist cognitive task relies on the self-organized dynamics of the whole system. In the traditional cognitivist approach however, a similar mechanism would be modeled in terms of explicit representations and memory.

### 3 The dynamics of anticipation

In the current paper we intend to present our work towards anticipatory cognitive architectures, with an emphasis on the role of non-neural internal states. Thus far, we have discussed a fundamental premise demonstrating how the dynamics of the body, and in particular its bio-regulatory processes, might be partially constitutive of cognition.

Operative definitions of anticipatory behavior stress the effect that an estimation of the future state of the system has on its current behavior [25,26]. Anticipation endows a cognitive agent with the capacity for faster and smoother action execution, facilitate action initiation, improve information seeking, decision making, predictive attention and social interaction [27,28,29,30]. In a recent paper, Butz argues how an anticipatory tension might influence both the development of neural structures and bias the agent to anticipatory behavior [31].

We suggest that settling on a dynamic attractor (e.g. see [24]) constitutes an implicit form of anticipation in at least one important sense. Once engaged with an attractor, the system enters a stable and fully determined regime. Our capacity to predict the trajectory of a strange attractor might well be limited by the confidence in the prediction that we can draw, as the system's non-linearities amplify our error. Nevertheless, once settled on an attractor that currently satisfies specific functional requirements, the whole dynamic of the system is attuned to a specific flow of events. An example of this attunement and its anticipatory role is Pavlov's dog, that salivates when food is made potentially available, thus effectively preparing its body for the digestive process. Some authors consider this kind of anticipation so important for an agent that conditioning, the prototypical basic form of learning in organisms, can be interpreted as mainly functional to its potentiation for originally neutral stimuli become suitable for the elicitation of anticipatory responses [32].

This observation constitutes a second important premise for what follows. To summarize, the body (in an extended sense that includes its non-neural internal mechanisms) constitutes a critical component of the potential dynamical richness of an agent attuned to its environment. Such richness, when autonomously viable, is intrinsically endowed with anticipatory power.

### 4 The bodily path of anticipation

A brief example of a prototypical situation will shed some light on our proposal. Let us consider a cognitive agent engaged in some activity, for example light-heartedly roaming on a soft lawn, enjoying the sight of colorful flowers and picking up wild berries. Suddenly something unexpected pops up from the bushes,

something potentially noxious and maybe never experienced before (e.g., depending on the agent’s particular sensitivity it could be a spider, a coral snake, or even a carnivorous dinosaur). We can be quite confident that a viable evolved agent would find a way to inhibit or redirect its current activity towards a more conservative behavior.

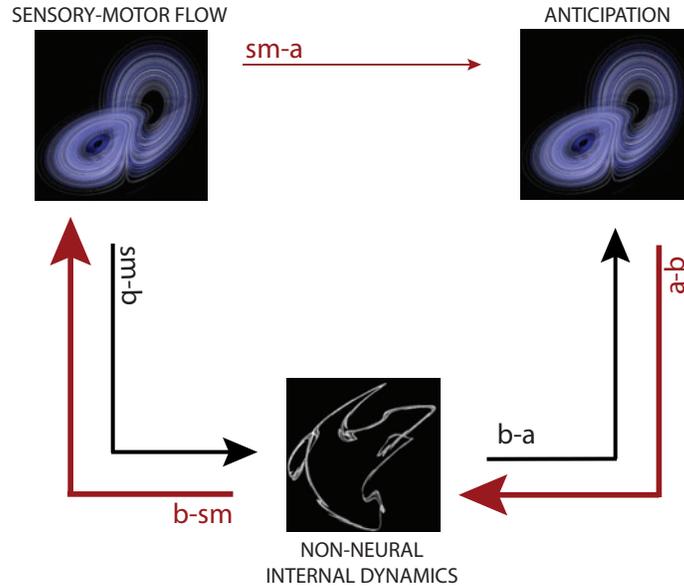
With reference to Fig. 3, we have to state a few preliminary assumptions:

- the global activity determining the current behavioral engagement between the agent and its environment (namely, a behavioral attractor, similar to [22,33,34]) is described by a few global variables that compress the specific relevant information for the current sensory-motor activity out of the enormous number of degrees of freedom of the system [12];
- the box labeled SENSORY-MOTOR FLOW represents the neural activity associated with the sensory-motor flow of the behavioral attractor;
- the corresponding non-neural bodily dynamics are summarized by the box labeled NON-NEURAL INTERNAL DYNAMICS;
- embedded within the global current dynamic we recognize a neural sensory-motor emulator<sup>3</sup> (box ANTICIPATION), whose evolution over time is dynamically correlated with the actual sensory-motor flow (similar to [33]), although not necessarily identical to it (as in [35]);
- the dynamics of the emulator (adapted during the evolutionary history and/or during the agent’s ontogenesis) can anticipate, in the dynamic sense illustrated above, the sensory-motor consequences of the engagement with a potentially noxious activity, as they follow a faster time scale.

Crucially, the capacity to predict the potential negative outcome endows the agent with a massive advantage: it attains the possibility to prepare itself before confronting the consequences of its current behavior, or to inhibit its behavior altogether. If we assume the possibility of a direct interaction between anticipatory and actual sensory-motor dynamics (i.e. a direct path between the boxes ANTICIPATION and SENSORY-MOTOR FLOW in Fig. 3), we immediately recognize a critical problem. Which kind of dynamics would eventually emerge after the current action is inhibited? Obviously, the dynamic structure emerging in the emulator should elicit a viable alternative behavioral attractor. How would that be selected?

Generalizing our example to other situations critical for the agent’s viability, our *bodily-anticipation hypothesis* is that, rather than a direct influence on the current behavior, the effect of the prediction of the emulator is actually mediated by the body. The outcome of the emulator affects the actual bodily dynamics (path a-b in Fig. 3), and altered bodily quantities transiently act as control parameters for the actual sensory-motor flow (path b-sm). Hence, the problem of the determination of the next behavioral attractor is off-loaded onto the bio-regulatory dynamics of the body. Destabilized by the input from the

<sup>3</sup> We follow here the terminology introduced by Grush [28], to denote an explicit subsystem that dynamically generates a prediction of the agent-environment sensory-motor interaction.



**Fig. 3.** Illustration of the bodily-anticipation hypothesis. In its roaming, our agent gets engaged with a potentially noxious interaction. Neural sensory-motor anticipatory dynamics, here conveniently isolated within the global coupled system (box labeled ANTICIPATION), predict the risk by determining a change in the current non-neural bodily dynamics (box NON-NEURAL INTERNAL DYNAMICS) through path a-b. This induces the agent to a visceral reaction, *as-if* actually engaged in the noxious sensory-motor experience. From here, indirectly through a further path b-sm, the anticipatory dynamics modulate the actual sensory-motor dynamics (box SENSORY-MOTOR FLOW). Following a quick reorganization of its behavioral attractor, our agent is attuned to escape the danger thanks to the mediation of its body, as there is no direct neural coupling path between anticipatory and sensory-motor dynamics.

sensory-motor emulator, the body viscerally reacts *as-if* actually engaged in such sensory-motor experience, eliciting behaviors that pull back the system towards viable regions. That implies that the body can achieve homeostatic balance not only in virtue of isolated non-neural internal dynamics, but also by triggering the selection of an appropriate behavior (path sm-b). This mechanism exploits the knowledge 'accumulated' by the body during a long and complex process of evolutionary and ontogenetic adaptation, functional to the viability of the agent. Equivalent knowledge, in case of a (theoretically possible) neural path directly coupling anticipatory and sensory-motor dynamics (through the missing path a-sm), should be somehow achieved by the nervous system.

## 5 The bodily path hypothesis put to the test

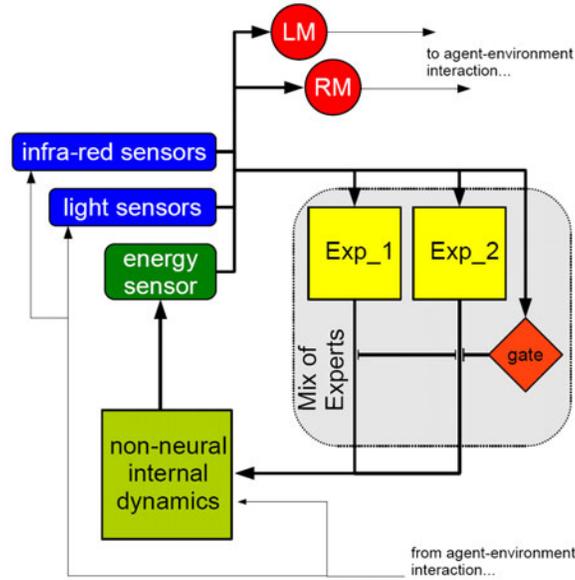
### 5.1 Implementation

The present section describes the first experimental steps toward a minimal implementation of the architectural plan outlined in Section 4. The experimental task takes place in the same simulated square arena and with the same agents as described in Section 2. The experimental task is extended to a scenario that can be abstractly likened to a *go-no go* task (loosely inspired by e.g. [36]). The light sources in our simulated setup emit according to two different patterns<sup>4</sup>. The first matches exactly the characteristics used in the experiment described above (continuous sensory regime), i.e. each light source emits a continuous steady level of luminance. The sensory consequences for the agents have already been demonstrated in Figure 2. Under this regime, nothing differs in the task with respect to the experiment described in Section 2. The agent, whose energy level is subject to a linear time decay (-0.008 per time step), is rewarded with an instantaneous full energy recharge upon invasion of the recharging area. The second sensory pattern is different in that the sensory input is rhythmically set to zero every third time step (intermittent sensory regime). This implies that during this modality the agent is subject to regular intervals of blindness. As pointed out elsewhere [22], the agent evolved in the previous experiment is robust enough to cope with massive perturbations, even of this nature, with no significant alteration of its behavior. Under intermittent regime, entering in the recharging area determines a punishment (-0.08 per time step) that speeds up the linear time decay. Each individual, whose lifetime lasts 1200 time steps, experiences the continuous sensory regime during time intervals [1, 200], [501, 700] and [1000, 1200]; intermittence occurs in the intervals [201, 500] and [701, 1000]. Severe punishment (-1000) was integrated in the fitness score in case of crashing against the walls.

A neurocontroller, assembled as a simple implementation of the general architecture introduced in Section 4, is sketched in Figure 4. We deployed simple feedforward artificial neural networks with no hidden layers, extracted from the population of the best individual in the previous experiment. Each ANN works in parallel with a Mixture of Recurrent Experts [37,33,38] whose role is to discriminate between the two different sensory regimes (continuous and intermittent). As each expert tries to outperform the other by generating the most accurate prediction on the sensor’s activity at the next time steps, they are actually chunking the sensory-motor flow according to its basic dynamical characteristics, as illustrated in [33].

In other words, each expert, by suppressing the output of the other, signals the engagement of the system with a specific sensory-motor flow, i.e. a specific stream of coupled perception and action. The gating sequence of the mixture of experts (that is, the description of which expert is currently active) can be mapped onto a binary variable. When the continuous sensory regime is detected,

<sup>4</sup> In a more natural metaphor, this might model the case of a succulent berry whose external pigmentation is different when unripe (and toxic) or ripe (and energizing).



**Fig. 4.** Sketch of a minimal implementation of our anticipatory cognitive architecture. Infra-red, light and energy sensors drive the two motorneurons through a feedforward ANN with no hidden layers. They also constitute a sensory flow that is processed by a mixture of recurrent experts. Each expert specializes on a specific sensory regime, and the gating signal perturbrates the non-neural internal dynamics of the agent.

nothing differs with respect to the dynamics used in the previously described experimental scenario. On the other hand, during the intermittent sensory regime, the energy level mechanism is overridden by a different mechanism, where the decay rate is freely evolved under the conditions specified above.

## 5.2 Results

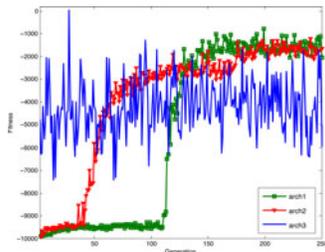
A standard evolutionary algorithm was run in this new scenario in order to select appropriate parameters for the neurocontrollers. Each agent was tested on its capacity to maximize the integral of its energy level over its lifetime (averaged over 10 epochs per individual), starting from random positions in the square environment.<sup>5</sup> Nevertheless, during continuous regime, stationary behaviors within the recharging area are discouraged, as the energy level is not integrated in the fitness until the agent leaves the rewarding space under this condition.

In this paper we compare the results for three different neurocontrollers:

<sup>5</sup> In order to partly make up for more advantageous starting positions, the first 100 time steps were non computed in the fitness function.

1. The basic feedforward architecture described in Section 2, whose weights and biases have been evolved from scratch on the new task.
2. As above, with evolution starting from the population of the best individual resulting from the previous experiment.
3. The minimal implementation of the general anticipatory cognitive architecture just introduced. The decay rate of the overriding mechanism for the energy level during intermittent regime is the only parameter modified by the evolutionary algorithm, as the rest of the networks remain frozen.

Figure 5 reports the fitness of the best individual (averaged over ten epochs of 1200 time steps each) for the best replication of the experiments for each of the three architectures described above (the parameters used in the evolutionary algorithm are shown in Table 1). Evolving weights and biases for the whole feedforward architectures (arch1 and arch2) produces similar results in terms of final performances. Nevertheless, the evolutionary process is much quicker when it can develop on the basis of the best population evolved on the simplified task presented in Section 2, although there is no initial advantage in this condition (the two curves basically overlap during the first 40 generations). The evolution of the new architecture (arch3) produces the best absolute performance and the bootstrapping of its performance is immediate. The evolutionary algorithm tends to select values for the single evolved parameter so that the energy level sensed during the intermittent regime simulates high energy. Therefore, consistently with the previous experiment, a tendency towards a photophobic behavior is triggered. Nevertheless, the fitness curve in this condition shows a very high variance, and in the long run the best individuals of the other architectures tend to outperform it.



**Fig. 5.** Performances of the best individuals of the three neurocontrollers during the evolutionary process. Although facing a problematic exploration in its parameter space, the minimal anticipatory cognitive architecture (arch3, continuous line) achieves satisfactory performance without any bootstrap phase.

A qualitative behavioral analysis emphasizes the different strategies deployed by the two classes of architecture, the simple feedforward ANNs with no hidden layers on one hand and the anticipatory architecture on the other. In the experiment described in Section 2 we observed two main classes of strategy. The

first, briefly described above and extensively reported in previous work [22,23], is highly dynamic and determines, under the modulation of the non-neural internal control parameter, an overt engagement with a specific object of interest, i.e. the rewarding light source (dynamic engagement). A second strategy relies on the geometrical constraints of the environment: the agent draws ad hoc spatial trajectories in order to achieve the appropriate timing required for the task (stereotypical engagement). Interestingly, the two classes of architecture specialize in producing the two different strategies.

Figure 6 demonstrates the typical behaviors of fit individuals in the two classes. Architectures belonging to the first class (arch 1 and arch2) tend to produce stereotypical attractors. Under the modulation of the different regimes, tighter loops invading the rewarding area will collect frequent rewards during continuous sensory regime, whereas slightly wider loops will stay clear of the recharging area in order to escape the punishment during intermittence. Therefore, the agent ignores the local effect of the rewarding area on its energy level. These behaviors depend on spatio-temporal constraints, as changes in the geometrical characteristics of the environment or in the timing of the different regimes might induce a dramatic drop in terms of performance. Our minimal anticipatory architectures, on the other hand, tend to develop a dynamic engagement with the light source (i.e. moving towards it - continuous trajectory) during the continuous sensory regime and a similarly straightforward disengagement during the intermittent regime (moving to safe distance from the punishing light - dashed line). In our viable anticipatory architectures, photophobic and phototactic behaviors are constantly balanced in order to take the agent either sufficiently close to, or far from, the recharging area, according to the current sensory regime.

---

Table 1: Evolutionary Parameters

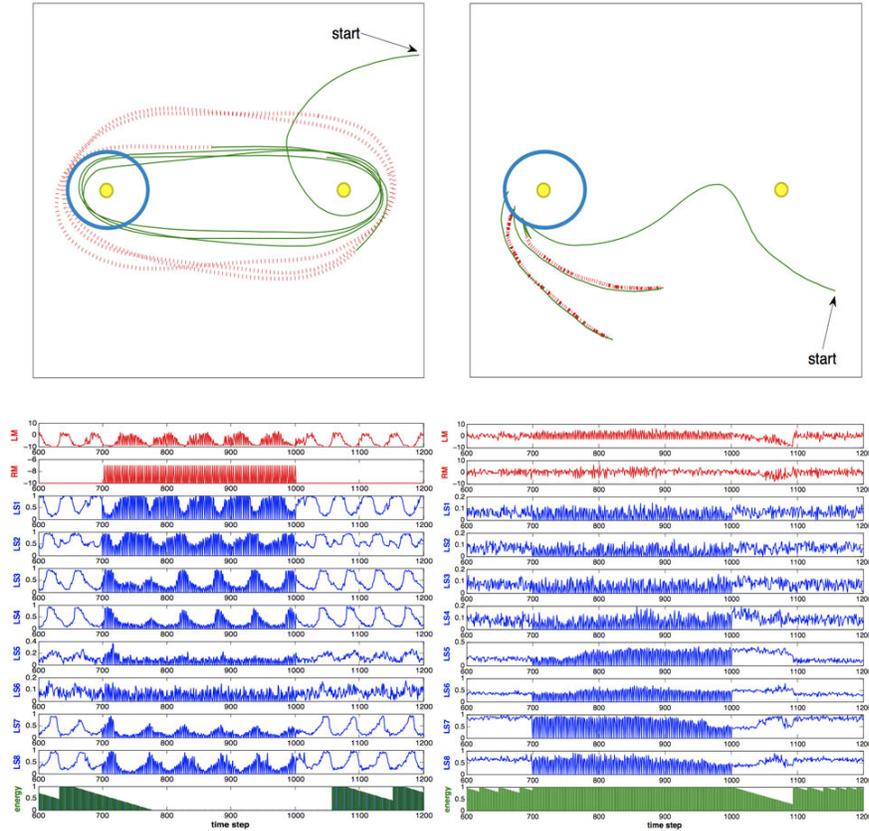
---

Number of generations:	1000
Number of individuals/generation:	100
Number of test epochs/generation:	10
Duration of one epoch (time steps):	1200 (= 120 s)
Starting position:	random
Probability of mutation:	0.02
Probability of crossover:	0
Sensory noise:	0.05

---

## 6 Discussion

The original task described in Section 2 (a stationary recharging area located in the proximity of a light source) and the extended task (alternate regime of reward and punishment on the same area) are obviously related. Nevertheless, it is interesting to notice that although even simple feedforward ANNs with no hidden layers can cope quite effectively with the new task, they achieve their skills after several generations of evolutionary adaptation. We obtain slightly better results starting the adaptation process from a population that already masters the orig-



**Fig. 6.** Typical spatial trajectories developed by the different architectures during evolutionary adaptation. **top left** - Simple feedforward ANNs tend to deploy a stereotypical strategy, i.e. their trajectories systematically engage in exploratory loops between the two light sources, entering the recharging area during the continuous regime (continuous line) and avoiding it during the intermittent regime (dashed line). **top right** - On the other hand, the behaviors that tend to emerge from our minimal anticipatory architectures dynamically engage and disengage with the rewarding/punishing area according to the different sensory regime (continuous/dashed lines represent the trajectories during continuous/intermittent sensory regimes). For a better resolution of details in the trajectory, the two pictures zoom on the area of main interest surrounding the light sources. **bottom** - The left and right panels exemplify, respectively for a feedforward and an anticipatory architecture, the activation of the two motoneurons (LM and RM), of the light sensors (LS1-8) and of the energy level sensor, during 600 time steps that include a double regime transition (continuous-intermittent-continuous) occurring at time steps 700 and 1000.

inal task. Nevertheless, in both cases we observe a slow bootstrapping, beginning with remarkably low performance (see Figure 5, arch1 and arch2). On the other hand, the minimal anticipatory architecture introduced in this paper, starting from the same population as used for arch2, demonstrates the instantaneous capacity to achieve satisfactory performance. Interestingly, the viable emerging behaviors with the anticipatory architecture are typically characterized by dynamical engagement and disengagement from the light sources, according to the current sensory regime. This results in flexible and robust behaviors, that contrast with the stereotypical behaviors achieved during the evolution of simple feedforward networks.

It might be observed that high variability of the data plotted in Figure 5 for our anticipatory architecture suggests a problematic evolution of its single free parameter. This is not surprising, since we maintained the exact same parameters for the evolutionary algorithm for all three architectures, and did not pay any attention to an optimal tweaking in this particular condition. In fact, the evolutionary process performs only slightly better on the anticipatory architecture than a random search (result not shown in detail). Nevertheless the best evolved neurocontroller achieves a fairly high (albeit isolated) performance, and for that particular value of the parameter even the average fitness of the population rises to the level of the best individuals (result not shown in detail). Rather, what should be emphasized is that the performance is achieved in parallel with a drastic dimensional reduction of the search space, and in these conditions even a random search can produce a number of individuals that immediately exhibit satisfactory behaviors in a dynamic scenario never experienced before. This is reminiscent of Ashby’s proposal of a ‘dumb’ mechanism that, in the need of maintaining a homeostatic balance for a set of *essential variables* critical to the agent’s survival, produce adaptive behaviors [8].

Obviously, our bodily-anticipation hypothesis does not rule out the feasibility of a totally disembodied and direct influence of the sensory-motor emulation on sensory-motor flow (the missing path a-sm in Fig. 3). Nevertheless, our approach drastically reduces the complexity of the problem of synthesis and adaptation. The search for viable parameters in the (potentially) massive dimensionality of the system’s degrees of freedom is reduced to a search in the subspace of the bodily parameters (in this case, the mere energy level). In our preliminary experiment, the search of the appropriate decay rate that is necessary to cope with the new task proves an effortless procedure for the minimalist anticipatory architecture. On the other hand, readaptation to the new task is cumbersome when we evolve the whole set of weights and biases in the ANNs. In a more naturalistic perspective, a basic organism constituted of a body coupled with a simple nervous system learns to survive first, by deploying a set of elementary sensory-motor reflexes in order to establish a basic form of viable coupling with its environment. This involves the evolutionary and ontogenetic adaptation of the interaction paths b-sm and sm-b in Fig. 3. Then the agent adaptively extends its viability by governing predictions. Incidentally, this is in accord with the design

principle of *holistic reductionism* [39], where the cognitive capacity of a minimal realization of a whole and viable autonomous system is incrementally extended.

The adaptation of the emulator takes place on the basis of the sensory-motor information provided along paths sm-a and b-a; paths ontogenetically adaptable, whose role could be adaptively weighted during the agent's life (as in [28]) and that might be transiently wiped out as the emulator proves its ability to produce effective predictions. In this extreme situation the system would express the capacity for 'blind navigation', i.e. navigation achieved by relying on its own sensory-motor predictions rather than on actual sensory information [35]. Therefore, the boxes ANTICIPATION and SENSORY-MOTOR FLOW in Fig. 3 act as informationally semi-permeable subsystems. Their coupling, from the latter to the former, should be modulated in a context-dependent way. The opposite coupling, far from absent, is indirectly realized via the body, according to our bodily-anticipation hypothesis.

In Section 2, we have described our model using the intuitive metaphor of an energy level mechanism, thus evoking biologically plausible dynamics of food intake and metabolism. Nevertheless, our intentionally simple scenario aimed to facilitate the abstraction to general principles. Metaphor aside, the fundamental aspect to consider is the coupling of different dynamic systems characterized by time scales that differ by several orders of magnitude (in particular we refer to the dynamics of the sensory-motor and the energy level systems). The availability of the slower dynamic of the energy level is exploited during the evolutionary adaptation of the system. In fact, the neurocontroller receives input vectors which are organized as dynamically related events in a continuous sensory-motor flow (i.e. contexts with a similar, although continuously varying, level of energy). The outcome of the adaptation process allows the system to integrate information over time. Although the sensory-motor mapping as such is purely reactive, this is not valid for the motor-sensory mapping and thereby for the behavior of the system as a whole. On the basis of these observations, we formulate the *hypothesis* that the access to a collection of attuned dynamic sub-systems characterized by intrinsic dynamics at different time scales and the exploitation of such differences, constitutes a powerful mechanism of embodied cognition, widely operating at the different levels of organization of biological cognition. A mechanism providing the cognitive system with the capacity to structure information on events which are relevant to its survival, with no need for explicit representations, memory or consciousness.

The focus on the role of multiple time scales, thus remapping the interpretation of our system in more abstract terms, dissolves the problematic distinction between non-neural and neural dynamics. We are advocating a mechanism where intrinsic time scales, characterizing mechanical, chemical and electrical phenomena in the body, might be coherently integrated into the cognitive process [17]. The dynamical richness of non-neural bodily processes might support the characteristic time scales of regular sensory-motor dynamics. The interest for the role of multiple time scales is currently growing in the neuroscientific community (e.g. [40,41]) as well as in cognitive robotics (e.g. [42,34,43,44]). A parallel might

be drawn with other neurocomputational architectures that deploy rich potential dynamics at different time scales, like Echo State Networks and Liquid State Machines [45,46]. Nevertheless, in the case of our architecture the bodily dynamics that inspire the non-neural internal mechanism are homeostatically and evolutionary relevant, i.e. they have a crucial effect on the body of the agent and on its behavior, independent on whether or not any cognitive process makes use of them (e.g. see [47]). Reservoir dynamics in ESN and LSM, on the other hand, can be completely random and irrelevant, and they have no effect whatsoever unless they are actually read out.

Our own and related experimental evidence in cognitive robotics supports our assumptions on paths b-sm and sm-b in Fig. 3, as examined in the previous Section 2 (e.g. [22,48,49,34]). This is to say that bodily states can modulate cognitive dynamics (e.g., think of the effects of particular chemical substances injected in your body) and particular behaviors can critically affect our body (e.g. in eating disorders). The capacity of the brain to anticipate sensory-motor correlates (path sm-a) is also supported by experiments in cognitive robotics, as in [33,34,35], and object of neuroscientific investigation (e.g., see [50]). In addition, the effect of mental imagery on non-neural bodily states is also rooted in neuroscientific evidence of biological cognitive processes (e.g. see [20]). The same author inspired the seemingly arbitrary choice to implement an overriding energy mechanism that takes over during intermittent regime. False bodily information can sometimes substitute for the correspondent actual state, for example when a contingent urge induces us to ignore pain [21]. Damasio seems to bring forth a somewhat opposite hypothesis, as he advocates *as-if body loops* [20,21]. In Damasio's theory, the emotional machinery, deeply integrated in the homeostatic mechanisms, plays a crucial role even in the case of highly logical functions, as in decision making. During the process of decision making it continuously supports the mental activity (*body loop*). After multiple exposure the brain builds appropriate neural causal associations that completely obliterate bodily information from the process (*as-if body loops*). Nevertheless, Bechara reports experimental results suggesting that *as-if body loops* are more plausible during choices made in highly predictable conditions (choice under certainty). As the decision process takes place in less predictable scenarios (full uncertainty) the *body loop* mode of operation becomes prominent [51]. We find this observation in perfect agreement with the intuition deployed in our model.

The architecture sketched in Figure 3 evokes a dynamic complexity that is drastically simplified in our initial implementation. The balance between the two subsystems ANTICIPATION and SENSORY-MOTOR FLOW is of course of the most delicate nature. In fact, via their effect on the non-neural internal dynamics, each of the two systems might simultaneously try to drive the system towards different dynamic attractors. This apparent contradiction should not necessarily be interpreted as a flaw in our proposed architecture. The dynamic tension between two competing requests maintaining the system in a regime of metasta-

bility<sup>6</sup>, rather than (more traditionally) of stability, might also be exploited as a potential opportunity. Some authors (e.g. [12]) consider metastability the fundamental state for a complex dynamic system like the brain, for it allows flexible and fast engagement and disengagement with contingent environmental requirements and constraints. Kelso, for example, offers an inspiring dynamic image of biological brains [ibid., p. 26]:

The human brain is essentially a pattern-forming self-organized system governed by nonlinear dynamic laws. Rather than compute, our brain "dwells" (at least transiently) in metastable states: it is poised on the brink of instability where it can switch flexibly and quickly. By living near critically, the brain is able to anticipate the future, not simply react to the present. All this involves the new physics of self-organization in which, incidentally, no single level is any more or less important than any other.

This intriguing scenario deserves further experimental investigation. Along a related line, we advocate the intrinsic unity of the general anticipatory architecture sketched in Fig. 3. We graphically split a system, that is actually meaningful only as a whole, into three different compartments just for the sake of clarity. The system should be conceived as a unity, where no component has a dominant role over the others, consistent with the final statement in Kelso's quote. Incidentally, under this perspective the traditional dichotomy between controlled and controller should be re-considered, as the different subsystems, through their coupling, mutually influence and regulate each other. Such interactions are graphically represented by the arrows in Fig. 3. Nevertheless, we might be interested in how a specific unbalance in one of the subsystems influences the others, and accordingly, as observers, choose the most convenient perspective. This step is legitimate and often even necessary to illuminate our analysis, albeit it does not modify the unitary nature of the system. As much as we emphasize the constitutional unity of our system, crucially linked with its environment, a more traditional symbolic approach would draw a sharp distinction between the agent and its environment.

## 7 Conclusions and future work

Non-neural internal states, in virtue of their different time scales, prove powerful potential props in support of cognitive processes. With this paper we hope to have contributed in some measure to highlighting their potential role, both in synthetic cognitive systems and, by extension, in biological ones. Preliminarily, we showed how a non-neural internal state, crucially characterized by a time scale that is orders of magnitude slower than ordinary dynamics of the sensory-motor interactions, can modulate the activity of a simple neurocontroller. What

---

<sup>6</sup> The concept of *metastability* can be intuitively introduced as a dynamical situation where the system does not express stable states, but a mere tendency towards them [12].

we achieved is the implementation of a self-organized, dynamic action selection mechanism, effectively operating in a context dependent way. Then we showed how these characteristics can be exploited by a minimal anticipatory cognitive architecture, using an explicit model for dynamic anticipation that operates via bodily mediation (*bodily-anticipation hypothesis*). This allows the system to scale up to more complex tasks never experienced before, achieving flexible and robust behavior with minimal adaptive cost.

Clearly, our hypotheses presented in Sections 4 and 6 require more experimental investigation and validation, which is currently under development in our lab. The work presented in this paper is still in progress and far from maturity. In order to facilitate the analysis and the extraction of general principles, our starting point is the synthesis of simple systems. A first extension will be the deployment of the full dynamic of the general anticipatory architecture sketched in Figure 3. The implementation of a more realistic internal dynamic, inspired by natural or artificial metabolic systems such as *microbial fuel cells* [47], represents the necessary step in order to systematically assess the potential of the architecture that we present here.

### Acknowledgments

The authors thank Silvia Coradeschi and Serge Thill for their comments on an early version of this paper, and Malin Aktius, Anthony Morse, Pierre Philippe and Henrik Svensson. We also thank the three anonymous reviewers for their suggestions. This work has been supported by a European Commission grant to the project *Integrating Cognition, Emotion and Autonomy* (ICEA, [www.iceaproject.eu](http://www.iceaproject.eu) IST-027819,) as part of the *European Cognitive Systems* initiative.

### References

1. Varela, F.J., Thompson, E.T., Rosch, E.: *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press (1992)
2. Thelen, E., Smith, L.B.: *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press (1996)
3. Clark, A.: *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA (1997)
4. Chrisley, R., Ziemke, T.: Embodiment. In: *Encyclopedia of Cognitive Science*. McMillan, London (2002) 1102–1108
5. Ziemke, T., Zlatev, J., Frank, R.M., eds.: *Body, Language and Mind: Embodiment*. Volume 1. Mouton de Gruyter, Berlin/New York (2007)
6. Nolfi, S.: Power and limits of reactive agents. *Neurocomputing* **42**(1-4) (2002) 119–145
7. Wiener, N.: *Cybernetics, or Control and Communication in the Animal and the Machine*. MIT Press (1965)
8. Ashby, W.R.: *Design for a Brain: The Origin of Adaptive Behavior*. Chapman ‘&’ Hall, London (1952)
9. Köhler, W.: *Gestalt Psychology*. Liveright (1947)

10. Gibson, J.J.: *The Ecological Approach To Visual Perception*. Houghton Mifflin (1979)
11. Van Gelder, T.: The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* **21** (2000) 615–628
12. Kelso, J.A.S.: *Dynamic Patterns: The Self-organization of Brain and Behavior*. MIT Press, Cambridge, MA (1995)
13. Beer, R.D.: Parameter space structure of continuous-time recurrent neural networks. *Neural Computation* **18** (2006) 3009–3051
14. Beer, R.D.: Dynamical approaches to cognitive science. *Trends in Cognitive Sciences* **4**(3) (2000) 91–99
15. Chiel, H., Beer, R.D.: The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences* **20**(12) (1997) 553–557
16. Suzuki, M., Floreano, D.: Enactive robot vision. *Adaptive Behavior* **16** (2008) 122–128
17. Parisi, D.: Internal robotics. *Connection Science* **16**(4) (2004) 325–338
18. Ziemke, T.: On the role of emotion in biological and robotic autonomy. *BioSystems* **91** (2008) 401–408
19. Ziemke, T., Lowe, R.: On the role of emotion in embodied cognitive architectures: From organisms to robots. *Cognitive computation* (2009-accepted)
20. Damasio, A.: *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Harvest Books (2000)
21. Damasio, A.: *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. Harcourt (2003)
22. Montebelli, A., Herrera, C., Ziemke, T.: On cognition as dynamical coupling: An analysis of behavioral attractor dynamics. *Adaptive Behavior* **16**(2-3) (2008) 182–195
23. Montebelli, A., Herrera, C., Ziemke, T.: An analysis of behavioral attractor dynamics. In Almeida e Costa, F., ed.: *Advances in Artificial Life: Proceedings of the 9th European Conference on Artificial Life*, Berlin, Springer (2007) 213–222
24. Strogatz, S.H.: *Nonlinear Dynamics and Chaos*. Westview Press, Cambridge, MA (1994)
25. Rosen, R.: *Anticipatory Systems*. Pergamon Press (1985)
26. Butz, M.V., Sigaud, O., Gérard, P.: Anticipatory behavior: Exploiting knowledge about the future to improve current behavior. In Butz, M.V., Sigaud, O., Gérard, P., eds.: *Anticipatory Behavior in Adaptive Learning Systems: Foundations, Theories and Systems*. LNCS, Berlin, Springer-Verlag (2003) 1–10
27. Butz, M.V., Pezzulo, G.: Benefits of anticipation in cognitive agents. In Pezzulo, G., Butz, M.V., Castelfranchi, C., Falcone, R., eds.: *The Challenge of Anticipation: A Unifying Framework for the Analysis and Design of Artificial Cognitive Systems*. Springer Verlag (2008) 45–62
28. Grush, R.: The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences* **27** (2004) 377–442
29. Barsalou, L.W.: Perceptual symbol systems. *Behavioral and Brain Sciences* **22** (1999) 577–660
30. Barsalou, L.W.: Social embodiment. In Ross, B.H., ed.: *The Psychology of Learning and Motivation*. Academic Press (2003) 43–92
31. Butz, M.V.: How and why the brain lays the foundations for a conscious self. *Constructivist Foundations* **4**(1) (2008) 1–42
32. Parisi, D.: *Mente: i nuovi modelli della Vita Artificiale*. Il Mulino, Bologna (1999)

33. Tani, J., Nolfi, S.: Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks* **12**(7-8) (1999) 1131–1141
34. Ito, M., Noda, K., Hoshino, Y., Tani, J.: Dynamic and interactive generation of object handling behaviors by a small humanoid robot using a dynamic neural network model. *Neural Networks* **19**(3) (2006) 323–337
35. Ziemke, T., Hesslow, G., A., J.D.: Internal simulation of perception: a minimal neuro-robotic model. *Neurocomputing* **68** (2005) 85–104
36. Schoenbaum, G., Chiba, A.A., Gallagher, M.: Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neuroscience* **1**(2) (1998) 155–159
37. Haykin, S.: *Neural networks: a comprehensive foundation*. Prentice Hall (1999)
38. Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, Geoffrey, E.: Adaptive mixtures of local experts. *Neural Computation* **3**(1) (1991) 79–87
39. Morse, A., Lowe, R., Ziemke, T.: Towards an enactive cognitive architecture. In: *Proceedings of the 2008 International Conference on Cognitive Systems*. (2008)
40. Kiebel, S.J., Daunizeau, J., Friston, K.J.: A hierarchy of time-scales and the brain. *PLoS Computational Biology* **4**(11) (2008)
41. Fusi, S., Asaad, W.F., Miller, E.K., Wang, X.J.: A neural circuit model of flexible sensorimotor mapping: Learning and forgetting on multiple timescales. *Neuron* **54**(2) (2007) 319–333
42. Yamashita, Y., Tani, J.: Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS Computational Biology* **4**(11) (2008)
43. Paine, R.W., Tani, J.: How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior* **13**(3) (2005) 211–225
44. Maniadakis, M., Tani, J.: Dynamical systems account for meta-level cognition. In Asada, M., Hallam, J.C., Meyer, J.A., Tani, J., eds.: *From Animals to Animats 10: Proceedings of the 10th International Conference on Simulation of Adaptive Behavior, SAB 2008*, Springer (2008) 3111–320
45. Jaeger, H., Haas, H.: Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science* **304**(5667) (2004) 78–80
46. Maass, W., Natschlagel, T., Markram, H.: Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation* **14**(11) (2002) 2531–2560
47. Melhuish, C., Ieropoulos, I., Greenman, J., Horsfield, I.: Energetically autonomous robots: food for thought. *Autonomous Robots* **21** (2006) 187–198
48. Tani, J., Ito, M.: Self-organization of behavioral primitives as multiple attractor dynamics: A robot experiment. *IEEE Trans. on Systems, Man, and Cybernetics. Part B* **33**(4) (2003) 481–488
49. Tani, J.: Learning to generate articulated behavior through the bottom-up and the top-down interaction processes. *Neural Networks* **16** (2003) 11–23
50. Hesslow, G.: Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences* **6**(6) (2002) 242–247
51. Bechara, A.: The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage. *Brain and Cognition* (55) (2004) 30–40