

# Embodied anticipation for swift re-adaptation in neurocomputational cognitive architectures for robotic agents

Alberto Montebelli, Robert Lowe, Tom Ziemke  
(alberto.montebelli@his.se, robert.lowe@his.se, tom.ziemke@his.se)

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

## Abstract

The coupling between a body (in an extended sense that encompasses both neural and non-neural dynamics) and its environment is here conceived as a critical substrate for cognition. We propose and discuss the plan for a neurocomputational cognitive architecture for robotic agents, so far implemented in its minimalist form for supporting the behavior of a simple simulated agent. A non-neural internal bodily mechanism (crucially characterized by a time scale much slower than the normal sensory-motor interactions of the robot with its environment) extends the cognitive potential of a system composed of purely reactive parts with a dynamic action selection mechanism and the capacity to integrate information over time. The same non-neural mechanism is the foundation for a novel, minimalist anticipatory architecture, capable of swift re-adaptation to related yet novel tasks.

**Keywords:** cognitive robotics; embodied cognition; dynamic systems; neuromodulation; anticipation; multiple time scales.

## Introduction: towards a cognitive robotic rendition of emotions

A systemic approach to the study of cognition permeates the general spirit of *gestalt psychology* (Köhler, 1947) and *ecological psychology* (Gibson, 1979), as well as the work of influential psychologists and philosophers (e.g., Piaget, Vygotsky and Merleau-Ponty). It also appears, as a modeling abstraction, in the seminal work of early cybernetics (Ashby, 1952; Wiener, 1965). In its modern form, the idea that *the whole is more than (and qualitatively different from) the sum of its parts* received a sound mathematical formalization through the science of non-linear dynamic systems (e.g., see Bergé et al., 1984; Haken, 2004) and pragmatic validation through physics. It constitutes one of the core theoretical milestones of contemporary science and influenced cognitive science with a whole new scientific paradigm, namely the Dynamic Systems approach to the study of biological cognition (e.g., Van Gelder, 2000; Kelso, 1995; Thelen & Smith, 1996).

The critical revision of the roles of body and environment in the cognitive process (Froese & Ziemke, 2009) constitutes the fundamental idea behind our paper. The systemic view conceives body and environment of the cognitive agent as constitutive of a largely distributed cognitive process, backing the brain in its operation by constantly offering cognitive support and tools. Thus, the cognitive process is the result of the activity of the brain-body-environment triad, whose components, dynamically coupled, are equally necessary to the mental process (Kelso, 1995; Clark, 1997). The body can be interpreted as an enduring pre/post-processor of neural information (Chiel & Beer, 1997), and its interaction with the

environment stores a wealth of knowledge about the "how to" of a cognitive activity (Pfeifer & Bongard, 2007). Research in embodied and situated cognition has investigated in theoretical and experimental terms the role of the body and of the environment in the cognitive process (Varela et al., 1992; Ziemke et al., 2007; Clancey, 1997). In this light, cognitive robotics (i.e., the use of robots as models of embodied and situated cognition) is the perfect candidate for generating an experimentally grounded synthesis, as it forces us researchers to take very seriously the interplay among coupled bodies, control systems and environments (Ziemke & Lowe, 2009; Parisi, 2004).

Alongside the role of the body projected towards its environment, there is a less obvious, less visible and consequently often neglected internal dynamic component of the body. We are referring to the plethora of background bio-regulatory mechanisms, aimed at the maintenance of a viable metabolic balance necessary for the organism's survival. An increasing number of researchers investigate the potential cognitive role of this hidden dynamic. Antonio Damasio illustrates a view of cognition deeply rooted in a hierarchy of bodily processes and consistent with state-of-the-art neurological findings (Damasio, 2000, 2003). According to Damasio, emotions emerge from the complex hierarchy that constitutes the levels of *automated homeostatic regulation* - the basic evolutionarily determined organization for the maintenance of the living organism. Similar approaches constitute the core motivations of *somatic theories of emotions* (Prinz, 2004; Panksepp, 2005). Indeed, grounding emotions in physical/physiological (rather than mental) terms constitutes the entry point for their appealing robotic rendition. In a recent paper, Domenico Parisi points to the necessity of a deep investigation of the relation between the control system and what happens inside of the body (Parisi, 2004). The emphasis on bodily parameters affecting control processes can be traced back further to the cyberneticist W. Ross Ashby, who focused on the behavioral consequences of a set of *essential variables*, critical to the organism's survival (e.g., sugar concentration in the blood and body temperature). The organism's implicit need to restrict their range within viable limits determines the random creation of new adaptive behaviors (Ashby, 1952). Focusing on the cognitive implications of bio-regulatory processes might be a promising direction for scientific explorations in order to implement robots endowed with genuine autonomy, agency, intentionality and meaningful interaction with their environment (Ziemke, 2008; Lowe

et al., 2008). Indeed, internal robotics alone is not sufficient for modeling emotions. It requires the presence of emotionally competent stimuli that derive from the coupling of body and environment in an adaptive history of interactions.

As a matter of fact, all the above is in contrast to the traditional perspective on AI and cognitive science, i.e., the presumption that the description of the world in terms of related symbol structures and logical processing on such structures is the necessary and sufficient condition for *general intelligent action* by *appropriate* instances of physical systems (Newell, 1980). A concept mapped in cognitive robotics onto the linear *sense-plan-execute* scheme, and conceptually akin to the functional approach of computational neuroscience, focused on specific and decontextualized sub-domains.

## From bodily neuromodulation to bodily anticipation

In recent minimalist cognitive robotics experiments we tested two different experimental scenarios (for detail, see Montebelli et al., 2008, 2007, 2009). In both experiments a simulated Khepera robot was free to move in a square arena, where two identical light sources, centrally located in the environment, cast a stationary light gradient. An invisible recharging area was centered under one of the two lights, randomly selected for each replication. The robot received sensory information through its light and distance sensors and moved according to the activation of two wheels controlled by a simple sensory-motor map, i.e., a single-layer, feed-forward artificial neural network (ANN). It also sensed its simulated energy level (e.g., the level of a battery charge), subject to linear decay, from a maximum value down to zero. In both scenarios, the fitness function rewarded at each time step the maintenance of positive levels of energy. Each individual was tested on runs of constant duration, for several replications. At the end of each generation, the best individuals were selected for reproduction according to a standard evolutionary algorithm.

**Experiment 1:** The entering of the recharging area provided an instantaneous full energy recharge. The evolutionary algorithm evolved weights and biases of the ANN. Obviously, the evolved agents performed well on such an elementary task. The interesting part of our work came when, setting aside the evolutionary task, we selected the best individual and used its energy level as control parameter of the agent-environment system. We clamped the energy level to a fixed value for the whole duration of each replication, and systematically explored values from empty to full in the different replications. Consequently, we were able to map the behavioral repertoire of the evolved agent as a function of its energy level. We observed three main classes of behavioral attractors: *exploratory behaviors* (i.e., the agent engages in large loops from one light source to the other - attractor class 'A'), *local behaviors* (the agent's loops are closely bound to a single light source - class 'C') and *hybrid behaviors* (combining the characteristics of both exploratory and local attractors - class 'B'). The expression of these three behavioral attrac-

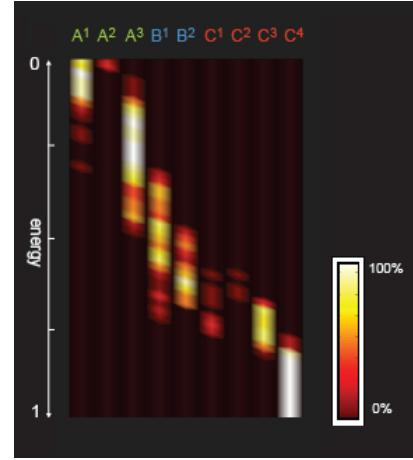


Figure 1: The intensity of the pixels for each column (corresponding to behavioral attractors belonging to classes A-C, as specified by their labels on the top row) represents the relative frequency of the attractor as a function of the energy level. Data from 500 replications (10 for each energy level). Adapted from (Montebelli et al., 2008)

tors was neatly distributed as a function of the energy level. Exploratory behaviors dominated the lowest range of energy levels, whereas local behaviors the highest ones (Figure 1). For intermediate levels of energy we found the prevalence of hybrid behaviors.

In sum, we showed how: 1) Minimalist non-neural bodily states (e.g., the energy level in our experiment) can *modulate* the sensory-motor map implemented by an ANN, and thus the behavior of the simulated robotic agent coupled with its environment. 2) This modulation can be exploited as a *dynamic action selection mechanism*. During the evolutionary task different classes of behavioral attractors were locally available to the agent, depending on its energy level. For example, an energy level of 0.7 (ref. Figure 1), led to the expression of attractor C3 (in 70% of the replications), C1 (20%) or B1 (10%). The actual selection of the specific attractor depended on the basin of attraction in which the combination of the starting position and the integrated effects of noise induced the system dynamics. 3) The cooperation between dynamics at different *time scales* can boost the cognitive potential of the system. In the case of our experiment (where the energy level mechanism was one order of magnitude slower than the normal sensory-motor interactions), a collection of purely reactive components was endowed with the capacity to integrate information over time (see Discussion).

**Experiment 2:** As before, a stationary gradient of environmental luminance (*continuous sensory regime*), correlated with a rewarding area centered on a randomly selected light source. However, during each replication this regime alternated with an *intermittent sensory regime*, where the light sources were obscured every third time step. Under this new condition, the randomly chosen area determined a punishment

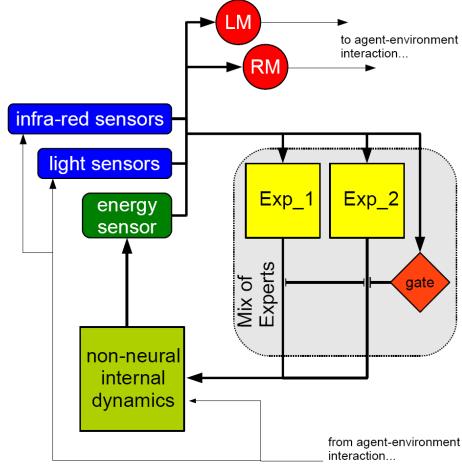


Figure 2: Minimalist anticipatory architecture. The sensory information (infra-red, light and energy sensors) drives the left and right motors (LM and RM) through a feedforward ANN with no hidden layers. The sensory flow is also processed by a mixture of recurrent experts, pre-adapted so that each expert is tuned to a specific sensory regime. The information on the current best expert (corresponding to one of the two regimes) is given by the gating signal, that selects the current energy mechanism of the agent.

ment in the form of an energy leak. As a biological metaphor, this alternation between regimes models the case of a succulent berry whose external pigmentation is different when unripe (and toxic) or ripe (and energizing). Again, the goal consisted in maintaining a positive energy level. We compared the simple architecture described in the previous experiment with a novel minimalist anticipatory architecture. In the former case, the evolutionary algorithm adapted the ANN's weights and biases on the new task, starting either from the final population evolved in the previous experiment or from a randomly generated population. In the case of the new architecture, shown in Figure 2, the original ANN (i.e., the simple ANN, whose weights and biases were extracted adopting the final population evolved during the previous experiment) was backed by a pre-adapted *mixture of recurrent experts* (Tani & Nolfi, 1999) that processed the sensory flow. During its adaptation, each expert competed with the others in order to generate the best prediction of the sensory state at the next time step. By doing so, two different experts became specialized by tuning to the specific dynamic flow of the two different regimes. Crucially, in the new architecture the activation of the expert tuned to the intermittent sensory regime triggered a new energy mechanism that overrode the original one. The decay rate of the overriding energy mechanism, rather than hardwired as before, is the one single parameter adapted by an evolutionary algorithm on the new task.

In short, we found that: 1) The systems provided with the

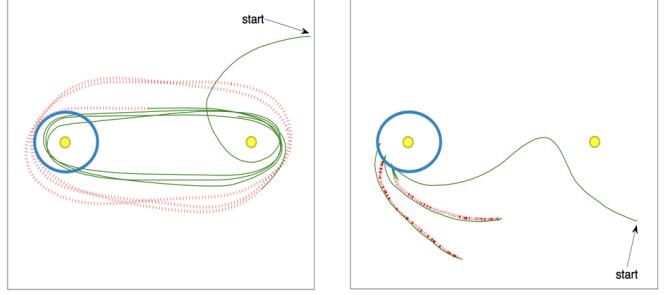


Figure 3: Prototypical spatial trajectories developed by the different architectures during evolutionary adaptation. **left:** Agents provided with feedforward ANNs tended to deploy a stereotypical strategy. Their trajectories systematically engaged in exploratory loops between the two light sources, entering the recharging area during the continuous regime (continuous line) and avoiding it during the intermittent regime (dashed line). **right:** Our minimal anticipatory architecture showed dynamical engagement/disengagement with the rewarding/punishing area according to the sensory regime.

anticipatory architecture developed an effective dynamic relation with its environment. They demonstrated a straightforward engagement with the rewarding light source during the continuous sensory regime, and a swift disengagement from the penalizing one during intermittent regime (ref. Figure 3, right). On the other hand, systems provided with the original ANN architectures tended to cope with the new task by relying on stereotypical behavioral attractors (Figure 3, left). During the continuous sensory regime they engaged in loops containing both light sources, approaching them close enough to enter their potential rewarding areas. During the intermittent regime they simply relaxed their trajectories with respect to the light sources, keeping at a slightly larger distance from them and consequently clear from the critical area, thus avoiding the punishment. This behavior ignores the effect of the recharging area on the energy level, merely relying on light sensor information and geometrical constraints. 2) In the case of the anticipatory architecture, the adaptive process for the new task proved easy, as even a random search could immediately generate agents with satisfactory performance. The evolutionary search was much more problematic for the original ANN, evolved from both starting conditions.

### A synthesis: the bodily-anticipation hypothesis

We will try to formalize the previous results in a general scheme. We have just seen how non-neural internal dynamics can modulate the current modality of the agent-environment interaction (i.e., its current behavioral attractor). On the other hand, the current behavior determines the current non-neural internal dynamics (e.g., an effective behavior that satisfies the experimental task maintains a high energy level). This bidirectional relation is expressed by the arrows connecting the

blocks labeled SENSORY-MOTOR FLOW and NON-NEURAL INTERNAL DYNAMICS in Figure 4. The former block represents the dynamic of the degrees of freedom relevant to the current sensory-motor engagement between the agent and its environment. Similarly, the latter embeds the relevant non-neural internal dynamics. In parallel, current sensory motor flow and internal dynamic drive a neural emulator block (labeled ANTICIPATION) that is capable, in virtue of its evolutionary history and/or ontogenetic adaptation, of dynamic anticipation. We suggested elsewhere (Montebelli et al., 2009) that a cognitive system settled on its behavioral attractor constitutes an important instance of an implicitly anticipatory system. In fact, the engagement with the attractor binds the system to a stable and qualitatively determined dynamic flow. An autonomous and viable dynamic is inherently endowed with anticipatory power. The main practical function of this emulator is to tune to the current sensory-motor dynamic and dynamically perturb the bodily dynamics with the anticipated consequences of the current dynamic interaction.

For example, consider a specimen agent, a caveman engaged in a relaxing and innocuous activity, e.g., picking berries in a forest. Out of the blue, an emotional stimulus, e.g., an apparently hungry, massive dinosaur, loudly enters the scene. The enormous time gap that separates the extinction of dinosaurs and the appearance of the first hominids is part of our example. We want to make sure that our specimen is experiencing a novel situation (therefore, a positivist caveman, who only brings solid scientific arguments to prove the dinosaur's anachronism, would be the perfect candidate for premature exhaustion of his own pedagogical role). The caveman's anticipatory system has no difficulty in predicting the most likely future scenario. The sensory-motor flow correspondent to the ongoing activity (picking berries) must be inhibited and redirected to a more conservative attitude. How will the next viable behavior (e.g., an impulsive fleeing) be selected? With this question in mind, our experiment explored the feasibility of a body-mediated pathway (arrow a-b in Figure 4). We tested the hypothesis that the anticipatory block (minimally implemented as the mixture of recurrent experts) might directly influence the non-neural bodily dynamics. In our prehistoric example, that means that once he perceived the emotional stimulus, our caveman would physically experience his own body torn by the fangs and nails of the dinosaur. It is likely that the caveman's evolutionary history and his ontogenesis had already created viable correlations between his dramatic visceral reaction and his fleeing for life, although the specific situation had never been experienced before. This constitutes the essence of our *bodily-anticipation hypothesis*: the selection of the next viable action is off-loaded onto the bio-regulatory dynamics of the body. Destabilized by the anticipated effect of the current interaction, the body reacts *as if* actually engaged in such sensory-motor experience. The bodily perturbation elicits reactions, already stored in the potential of bodily and neural interactions, that tend to pull the system back into viable regions.

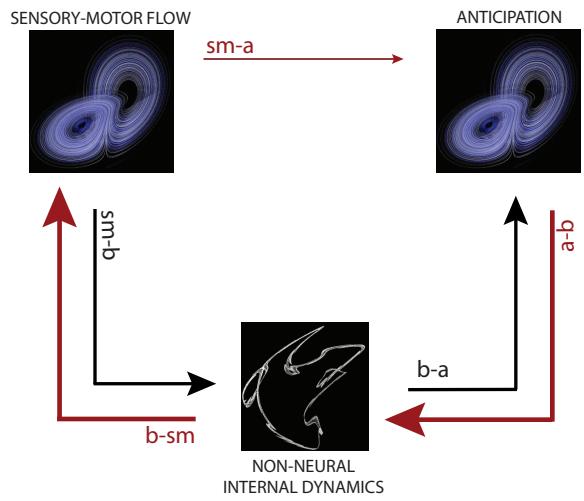


Figure 4: Illustration of the bodily-anticipation hypothesis. During its daily roaming, our agent gets engaged with a potentially noxious interaction. Neural sensory-motor anticipatory dynamics, here conveniently isolated within the global coupled system (box labeled ANTICIPATION), predict the risk by physically perturbing the current non-neural bodily dynamics (NON-NEURAL INTERNAL DYNAMICS) through path a-b and from there, indirectly through a further path b-sm, the actual sensory-motor dynamics (SENSORY-MOTOR FLOW). Following a quick reorganization of its behavioral attractor, our agent is attuned to face the novel danger thanks to the mediation of its body, without any direct influence of anticipation on the selection of the new behavior.

## Discussion

We hope to have clarified enough the importance of conceptualizing the phenomenon of cognition as emergent from the coupling of body, nervous system and environment. Within this systemic view, the boundary separating each subsystem is nothing but a useful artifice, functional to the analysis of a complex system dominated by circular relations. This tight coupling casts a light on an interesting matter. What is internal? What is external? We prefer to avoid such dichotomy, rather focusing on the system composed of dynamically interacting parts. At any given time its dynamic balance will be perturbed by stimuli coming from different sources (e.g. the environment, the agent's regulatory mechanisms, its nervous system). Each perturbation would produce a consonant reaction of the system's trajectory in its phase space. Each time, according to the needs of the analysis, we will have to properly redraw the boundary between input and output, cause and its effect. Parisi suggested objective criteria for partitioning the inside and outside of the body in natural agents, on the grounds of the physical-chemical processes that tend to dominate the two interfaces (Parisi, 2004). Local and specific interactions with fast dynamics, archetypal of physical

processes, tend to characterize the interface with the external world. Global and diffused variations with slower time scales, characteristic of chemical processes, tend to take place inside the organisms. The generalization to the realm of cognitive robotics prepares us for the next observation.

Could our non-neural internal dynamic be translated into purely neural mechanisms? We dissolve the problem in its abstract formalization. The interplay of the different time scales that characterize the energy mechanism and the other sensory-motor interactions with the environment is crucial to our model. During the artificial evolution of the system, the slower dynamic of the energy level organizes the continuous sensory-motor flow in dynamically related events. As we observed, this endowed a system composed of purely reactive elements with the capacity to integrate information over time. Elsewhere (Montebelli et al., 2009), we conjectured 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." With this in mind we can look at the plethora of bio-regulatory phenomena with new eyes. The characteristic time scales of non-neural bodily processes might provide exactly the kind of dynamical richness that we are advocating. The role of multiple time scales is currently attracting the attention of the scientific community, both in computational neuroscience and cognitive robotics (e.g., Kiebel et al., 2008; Yamashita & Tani, 2008; Paine & Tani, 2005).

The paths in the general scheme sketched in Figure 4 are less arbitrary than they might look at first glance. Our own and related work in cognitive robotics (Montebelli et al., 2008, 2007; Ito et al., 2006), motivates the paths sm-b and b-sm in Figure 4. The claim that the internal dynamics of the body (e.g., a sudden injection of adrenaline) affect the behavior and that behavior affects the body (e.g., eating or declining a fifth slice of your birthday cake) should not strike us as bizarre. The capacity of the brain to anticipate sensory-motor correlates (path sm-a) is currently the object of intensive research in neuroscience and cognitive robotics (e.g., Hesslow, 2002; Barsalou, 1999; Tani & Nolfi, 1999; Ito et al., 2006). Damasio (2000), reports the case of a professional musician who could systematically control her emotional machinery in experimental conditions (path a-b). Also the seemingly arbitrary switch between the natural energy dynamic and the overriding energy mechanism taking over during intermittent sensory regime is inspired by neurophysiological analogs. False bodily information can sometimes substitute for the actual state, for example, in the case of endogenously altered nociceptive signals. There is an obvious advantage for a wounded organism to ignore the pain when it is fleeing from

the danger that produced it (Damasio, 2003).

Obviously, our bodily-anticipation hypothesis does not rule out the possibility of a co-existence with a neural pathway between anticipation and sensory-motor flow (the missing path a-sm in Figure 4). Nevertheless, we point to the fact that our minimalist anticipatory architecture drastically simplifies the problem of readapting to a new task. The search space, characterized by the potentially enormous number of degrees of freedom of an ANN, is reduced by our bodily-anticipation hypothesis to the much smaller dimensionality of the bodily neuromodulators (the energy level in our minimalist example). An argument in favor of a mental path seems to be brought forth by Damasio, as he introduces the *as-if body loops* (Damasio, 2000). The emotional machine is in Damasio's theory central even to highly logical functions, e.g. decision making (Damasio, 2000). Its support can be elicited directly, but after repeated exposure the brain can build consistent causal associations and thus bypass the body in the decision process. Nevertheless, Bechara refers to preliminary results showing how in the process of decision making the role of the as-if body loop might be restricted to highly predictable situations. As the decision drifts from certainty to risk or ambiguity the *body loop* mode of operation, where the bodily mechanisms are directly engaged, becomes prominent (Bechara, 2004). We find this observation in perfect harmony with the intuition inspiring our model.

## Conclusions and future work

We showed experimental evidence of how non-neural internal dynamics, following a slow time scale, can modulate the activity of an ANN and consequently the behavior of an agent coupled with its environment. A traditional evolutionary algorithm self-organized this modulation, implementing a dynamic action selection mechanism. The coordination of multiple time-scales induced more sophisticated cognitive capacities in a very simple system. Furthermore, a novel minimalist anticipatory architecture provided flexible and dynamic engagement of the agent with its environment, as a swift adaptation to a brand new task was accomplished.

Our model might be accused of being an *ad hoc* arrangement: First, for the arbitrary decision to override the original non-neural internal mechanism (although the same strategy can be found in natural agents). Second, for selecting the decay rate of the overriding energy mechanism as critical parameter to be adapted by the evolutionary algorithm. This is a reasonable criticism. Nevertheless, given the extreme simplicity of our current setup, such design choices were necessary. We chose a minimalist model as a deliberate preference, for simplicity facilitates the focusing on general principles and detailed analysis. We think that such objections can be more easily confuted given a modification towards a slightly more complex model, both in terms of task and architecture. In particular, future work will specifically address the implementation of more realistic internal dynamics, inspired by natural metabolic systems. We believe that this work illus-

trates promising results in terms of basic organizational principles of cognition, that can be usefully explored by minimalist cognitive architectures. The new questions that it raises require and deserve further investigation and validation.

### Acknowledgments

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

Ashby, W. R. (1952). *Design for a brain*. London: Chapman & Hall.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.

Bechara, A. (2004). The role of emotion in decision-making. *Brain and Cognition*(55), 30-40.

Bergé, P., Pomeau, Y., & Vidal, C. (1984). *Order within chaos*. Wiley-Interscience.

Chiel, H., & Beer, R. D. (1997). The brain has a body. *Trends in Neurosciences*, 20(12), 553-557.

Clancey, W. J. (1997). *Situated cognition: on human knowledge and computer representations*. Cambridge University Press.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.

Damasio, A. (2000). *The feeling of what happens*. Harvest Books.

Damasio, A. (2003). *Looking for spinoza: Joy, sorrow, and the feeling brain*. Harcourt.

Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173, 466-500.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Haken, H. (2004). *Synergetics*. Springer.

Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences*, 6(6), 242-247.

Ito, M., Noda, K., Hoshino, Y., & Tani, J. (2006). Dynamic and interactive generation of object handling behaviors by a small humanoid robot using a dynamic neural network model. *Neural Networks*, 19(3), 323-337.

Kelso, J. A. S. (1995). *Dynamic patterns*. Cambridge, MA: MIT Press.

Kiebel, S. J., Daunizeau, J., & Friston, K. J. (2008). A hierarchy of time-scales and the brain. *PLoS Computational Biology*, 4(11).

Köhler, W. (1947). *Gestalt psychology*. Liveright.

Lowe, R., Herrera, C., Morse, A., & Ziemke, T. (2008). The embodied dynamics of emotion, appraisal and attention. In L. Paletta & E. Rome (Eds.), *Attention in cognitive systems* (p. 1-20). Berlin: Springer.

Montebelli, A., Herrera, C., & Ziemke, T. (2007). An analysis of behavioral attractor dynamics. In F. Almeida e Costa (Ed.), *Advances in artificial life: Proceedings of the 9th european conference on artificial life* (p. 213-222). Berlin: Springer.

Montebelli, A., Herrera, C., & Ziemke, T. (2008). On cognition as dynamical coupling: An analysis of behavioral attractor dynamics. *Adaptive Behavior*, 16(2-3), 182-195.

Montebelli, A., Lowe, R., & Ziemke, T. (2009). The cognitive body: from dynamic modulation to anticipation. In *Anticipatory behavior in adaptive learning systems: from sensorimotor to higher-level cognitive capabilities*. Springer.

Newell, A. (1980). Physical symbol systems. *Cognitive Science*, 4(2), 135-183.

Paine, R. W., & Tani, J. (2005). How hierarchical control self-organizes in artificial adaptive systems. *Adaptive Behavior*, 13(3), 211-225.

Panksepp, J. (2005). Affective consciousness. *Consciousness and Cognition*, 14, 30-80.

Parisi, D. (2004). Internal robotics. *Connection Science*, 16(4), 325-338.

Pfeifer, R., & Bongard, J. (2007). *How the body shapes the way we think*. Cambridge, MA: MIT Press.

Prinz, J. J. (2004). *Gut reactions*. Oxford University Press.

Tani, J., & Nolfi, S. (1999). Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12(7-8), 1131-1141.

Thelen, E., & Smith, L. B. (1996). *A dynamic systems approach to the development of cognition and action*. MIT Press.

Van Gelder, T. (2000). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615-628.

Varela, F. J., Thompson, E. T., & Rosch, E. (1992). *The embodied mind*. MIT Press.

Wiener, N. (1965). *Cybernetics, or control and communication in the animal and the machine*. MIT Press.

Yamashita, Y., & Tani, J. (2008). Emergence of functional hierarchy in a multiple timescale neural network model. *PLoS Computational Biology*, 4(11).

Ziemke, T. (2008). On the role of emotion in biological and robotic autonomy. *BioSystems*, 91, 401-408.

Ziemke, T., & Lowe, R. (2009). On the role of emotion in embodied cognitive architectures: From organisms to robots. *Cognitive computation*, 1(1), 104-117.

Ziemke, T., Zlatev, J., & Frank, R. M. (Eds.). (2007). *Body, language and mind: Embodiment* (Vol. 1). Berlin/New York: Mouton de Gruyter.