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

This is an author produced version of a paper published in Adaptive Behavior. This paper has been peer-reviewed but does not include the final publisher proof-corrections or journal pagination.

Citation for the published paper:

Montebelli, Alberto, Herrera, Carlos & Ziemke, Tom

"On Cognition as Dynamical Coupling: An Analysis of Behavioral Attractor Dynamics"

Adaptive Behavior, 2008 16(2-3): pp 182-195

URL: <http://dx.doi.org/10.1177/1059712308089180>

Access to the published version may require subscription.

Published with permission from: Sage Publications

On Cognition as Dynamical Coupling: An Analysis of Behavioral Attractor Dynamics

Alberto Montebelli, Carlos Herrera, Tom Ziemke
University of Skvde
School of Humanities and Informatics

Corresponding author:

Alberto Montebelli
University of Skvde
School of Humanities and Informatics
Hgskolevgen, PO Box 408, 54128 Skvde, Sweden
Email: alberto.montebelli@his.se
Phone: +46(0)500 448367 fax: +46(0)500 448399

Abstract

The interaction of brain, body and environment can result in complex behavior with rich dynamics even for relatively simple agents. Such dynamics are, however, often difficult to analyze. In this article we explore the case of a simple simulated robotic agent, equipped with a reactive neurocontroller and an energy level, that the agent has been evolved to re-charge. A dynamical systems analysis, shows that a non-neural internal state (energy level), despite its simplicity, dynamically modulates the behavioral attractors of the agent-environment system, such that the robot's behavioral repertoire is continually adapted to its current situation and energy level. What emerges is a dynamic, non-deterministic and highly self-organized action selection mechanism, originating from the dynamical coupling of four systems (non-neural internal states, neurocontroller, body and environment) operating at very different time scales.

Keywords:

evolutionary robotics, dynamical systems, embodied cognition, internal robotics, neuromodulation

Introduction and Background

In the dynamical systems (DS) approach to adaptive behavior and cognition (e.g. Kelso, 1995; Van Gelder, 2000; Clark, 1997) agents and their environments are viewed as tightly coupled dynamical systems. Nervous systems, bodies, and their environments are considered to have complementary roles in producing a rich range of adaptive behaviors (Beer, 1995; Chiel & Beer, 1997). Accordingly, cognitive agents are not closed systems whose activity can be reduced to the mapping of sensory inputs to motor outputs. In natural agents, nervous system, body and environment are not three independent components engaged in a synchronic interaction; rather, natural adaptation is the ongoing result of a global, self-organizing process. The implementation of artificial adaptation on the basis of principles of self-organization is therefore promising, and has been explored in much work in artificial life and adaptive robotics in recent years. The DS approach offers a range of tools available to the theorist aiming to understand adaptive processes as they unfold in time. In recent years there also has been a decisive turn towards aspects of situatedness and embodiment in much of cognitive science and AI (Clark, 1997; Pfeifer, 1999; Pfeifer & Scheier, 1999). Much recent robotics research has therefore focused on the sensorimotor interaction between control systems, their robotic bodies and their external environment. For example, it has been demonstrated that control may be highly distributed and, for instance, morphology can play an active role in control (Pfeifer, 1999).

Furthermore, Parisi (Parisi, 2004) has recently argued that “to understand the behavior of organisms more adequately we also need to reproduce in robots the inside of the body of organisms and to study the interactions of the robot’s control system with what is inside the body”. Parisi coined the term *internal robotics* to denote the interaction between the (neural) control system and the rest of the body. One important aspect of this interaction is motivation. Ashby (Ashby, 1957) used the term “homeostasis” to refer to the fact that certain variables, essential for the system, need to remain within certain limits – a “viability zone” (Meyer, 1995; Di Paolo, 2003) – for the organism to function and stay alive. That means, an adaptive agent needs to incorporate mechanisms that ensure that deviations from that viability zone are balanced during adaptive interaction. The viability zone is therefore ensured, not only through internal homeostatic mechanisms, but also motivating adaptive behavior.

There have been several attempts to model internal states as motivations for the control of robots (Gadanhó & Hallam, 2001; Avila-García & Cañamero, 2004). The role of internal states as motivations in artificial agents will strongly depend on what particular architecture the artificial agent embodies. If different forms of behavior are determined by functional modules, the specific values of internal variables can effectively influence an action selection process (Maes, 1991; D. Cañamero, 1997; L. Cañamero, 2003). Research on self-organizing behavioral systems has demonstrated that breaking down control systems into modules, according to some designer’s or observer’s preconception or perception of what ‘behaviors’ might be required, is not unproblematic (e.g. Nolfi, 1997; Ziemke, 2000). For example, it has been shown that behavioral adaptivity through quasi-modularity at a functional level can be realized without actual modules, through relatively simple recurrent neurocontrollers that make use of feedback to dynamically modulate their own behavior over time (Ziemke, 2000). Similarly, in recent work on recurrent architectures with parametric

bias values (Tani, 2003; Tani & Ito, 2002; Ito & Tani, 2004; Ito, Noda, Hoshino, & Tani, 2006) internal states function as control parameters of nonlinear dynamical systems to generate the dynamical bifurcation to distinct behaviors. In this case, the dynamics of internal states (the parametric bias) is achieved through anticipation and back-propagation of the error.

In this article, we further investigate the role of internal states in the modulation of adaptive behavior, exploring the hypothesis that even very simple internal states can modulate the functioning of a simple, non-specialist network to produce different behaviors. Our perspective is radically different from related work on action selection since we do not consider behaviors generated by a collection of structural modules. In other words, rather than focusing on how a selection mechanism can arbitrate the activation of a module that encapsulates some behavioral program, we investigate how internal states can modulate the activity of a control system so to generate distinguishable interactive dynamics and behaviors in different contexts. Using a novel dynamical systems analysis, we distinguish several attractors, which can be interpreted as behavioral modes. The dynamical mechanism of the internal states used in the experiments documented here (energy level), is intentionally kept very simple. The underlying hypothesis examined in the following analysis, is that, in a self-organized homeostatic system, internal states can play the role of control parameters of the interaction (Kelso, 1995), therefore affecting behavior not through the selection of subsystems/modules, but through modulation of the phase space that dynamically adapts the ‘behavioral repertoire’ locally available to the agent.

Method

In this article, we report an experiment based on an intentionally simple evolutionary robotics simulation model of dynamical energy management, inspired by the theoretical framework sketched above. The control system for a simulated Khepera robot, implemented by a fully connected feedforward artificial neural network with no hidden layer, was artificially evolved by a standard evolutionary algorithm within a square environment containing two identical light sources (see Figure 1). The agent receives inputs from eight infrared sensors and eight light sensors (in both cases, six placed frontally according to a bilateral symmetry and two in the back of the robot) and one energy sensor (tracking the current level of energy). All the inputs are directly connected, through the evolved weights, to the two motor neurons respectively driving the activation of the left and right motor (positive and negative values determine the rotation in the two possible directions). Two light sources are symmetrically located at the center of the arena (700 mm x 700 mm), their distance from each other being comparable to the maximum distance range of the light sensors (300 mm).

The robot was initially provided with an energy level of 1.0 (“full”). Its energy was subject to a continuous, linear decay at a rate of 0.008 per time step, down to zero. When the robot (that can be geometrically described as a cylinder of 50 mm diameter) entered a target area surrounding one of the lights (a circle of 30 mm radius highlighted, for the reader’s benefit, by the dotted circle in Figure 1) its energy reservoir was instantaneously and completely refilled, thus generating an internal effect sensed by the agent¹.

The fitness, integrated over the whole lifetime, is calculated as follows: a positive value of energy would result in a positive additive constant (+0.01 per time step), if the agent was positioned outside the target area (i.e. stationary behaviors inside the target area

Evolutionary Parameters	
Number of generations:	650
Number of individuals/generation:	20
Number of test epochs/generation:	18
Duration of one epoch (time steps):	600 (= 60 s)
Starting position:	random
Probability of mutation:	0.02
Probability of crossover:	0
Sensory noise:	0.05

Table 1: List of the parameters used during the evolutionary process.

are not rewarded). In order to facilitate the evolutionary bootstrapping, the performance of the first 200 time steps of lifetime was not integrated into the fitness. On the other hand, a null energy level would produce the opposite effect of integrating a negative constant to the fitness (-0.02). The survival of the individual was only limited by the number of time steps during its lifetime, which immediately ended in the case of eventual crashes against the walls, which also determined a severe punishment (-100). The simulated evolutionary process took place according to the parameters specified in Table 1.

Results and Analysis

At the end of the evolutionary process the simulated agent (endowed with the fittest evolved controller) proves reliably effective in disambiguating between the two light sources, according to their capacity to affect, neutrally or positively, its energy level. As exemplified in Figure 1, the agent, to whom the two light sources look identical, approaches a first light source and, receiving no reward (that is, sensing no increase in its energy level), quickly migrates towards the other light source that supplies it with energy. This determines the engagement of the agent within a stable behavior in the neighborhood of the rewarding light source.

The development of a skillful controller capable of light disambiguation should be only considered of minor interest and here prototypical for demonstration. Rather than the goal, it constitutes the actual starting point for our analysis, which will emphasize the general holistic nature of the relation between agents and their environments, the complex dynamics emerging from such interaction and the crucial role of non-neural internal states in influencing the emergent dynamics. During the evolutionary process, the energy level plays the generic role of one of the many internal state variables of the system. By ‘system’, we here mean the dynamical system constituted by the agent, its control system, its environment/arena and the set of bio-regulatory mechanisms (synthesized, as a crude first approximation, in the agent’s energy level). In the following analysis, where the evolved agent is tested under conditions that transfigure the evolutionary task, the energy level is used as control parameter of the overall system. The systematic manipulation of the energy

Figure 1. Evolved behavior. As its energy level linearly decays, the simulated agent approaches the light (represented by a filled circle) to the right of the environment. The lack of any effect on its energy level results in the following approach of the next light. As its energy reservoir is instantaneously refilled to the maximum level of energy, 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 current activation of the left and right motors; LS1-8 represent the activation of the light sensors in time; EnS displays the energy level.

level allows us to visit a wealth of behavioral attractors embedded in the system during its evolution (where the adaptation is limited to the reactive neural controller).

Energy binary-switch

In order to lay down the most extreme possible scenarios, we ran some simulations where the same agent as illustrated in Figure 1 was placed in one of a few possible initial positions and fed with a steady high level of energy. After a suitable interval of time, chosen to exhaust the transitory phase from the initial position to the development of the specific behavioral attractor, the energy level was suddenly switched and clamped to zero. Consequently, none of the two light sources would any longer provide the agent with an energy recharge.

As shown in Figure 2, during the first 200 time steps the agent drifts at low speed away from the sensed light sources, moving from its initial position. As the energy level switches from 1.0 down to 0.0 (illustrated by the abrupt switch in the graph of the energy level) the magnitude of the agent's speed is abruptly increased, as clearly shown by the sudden increment of its motor activation (the negative values indicate that the robot is moving backwards). The agent engages in a fast run, alternating visits to both of the light sources.

Energy clamp

The interesting behavioral dynamics observed in the previous experiment encouraged the development of our analysis, accomplished by refining the quantization of the possible energy levels, as inspired by voltage clamp techniques, routinely used in neurophysiology (e.g. Nicholls, Martin, Wallace, & Fuchs, 2001). Accordingly, the energy sensor input of

Figure 2. Energy binary-switch experiment. During the first 200 time steps (energy level 1.0; dashed trajectory), the same agent as illustrated in Figure 1 slowly moves forward (as the low positive motor activations demonstrate) from its initial position towards the closest light. As the activation of the frontal light sensors increases, the simultaneous change in the activation of the two motors triggers a sharp turn (around time step 50). Then, still in slow forward movement, the robot drifts away from the light source. When the energy level switches to zero (time step 200) its motor activations suddenly assume high negative values (backward high-speed motion) as the agent alternately visits the two light sources.

the neural controller was systematically clamped to discrete, steady levels of energy, ranging from 0.0 to 1.0, with constant increments of 0.02. For each energy level the agent was left free to interact with its environment, starting from 10 fixed positions and orientations, for a time sufficient to fully develop its ongoing dynamics in a wider environment (1500 mm x 1500 mm), broadened so to prevent any sensory interference from the walls during exploratory behaviors. The starting positions were chosen according to the specific geometrical symmetry of the environment and within a distance from the closest light source so that the agent's light sensors could actually detect it.

A taxonomy of behavioral attractors.

Following the qualitative observation of a large set of different trajectories in the phase space, we developed a taxonomy consisting of nine behavioral attractors², based on their general morphological and spectral characteristics. Their iconic representations are sketched in Figure 3, as projection of the phase space on the physical space of the robot's environment. In order to make sure that transients have already converged to stable trajectories, all the panels in the picture are based on the last 1000 time steps (100 simulated seconds) from replications lasting 3000 time steps each. An analogous attention to the full development of the behavioral transients characterizes also the rest of the following analysis.

Spectrum analysis by Fast Fourier Transform (e.g. in Haykin & Van Veen, 2003) proved particularly powerful in supporting the discrimination of visually similar behaviors. For example, attractors labeled F and H, although clearly different in Figure 3, might be hard to visually disambiguate over a number of repetitions (details not shown here). However, as illustrated in Figure 4, their spectra in the frequency domain in fact unmistakably

Figure 3. Iconic representation of the spatial trajectories for the nine behavioral attractors in our taxonomy (plotted over 1000 time steps after convergence of the transients). The scale has been systematically varied for the reader's sake, in order to facilitate a deeper qualitative grasp of the attractor's morphology (where possible, light sources are indicated by a star). Basic morphological elements in the trajectories are expressed by the different attractors in different combinations.

and reliably characterize them in a way that not only simplifies their classification, but also clearly demonstrates that they should be considered as two distinct behaviors. In fact, while the spectrum of attractor H shows sharp peaks forming a sequence of harmonics at frequencies multiple of 0.45 Hz, thus revealing its highly periodic nature, attractor F has a quite continuous spectrum, dominated by low frequency harmonics and a peak centered around 0.2 Hz.

Energy level as dynamical behavioral modulator.

The effect of the energy level as control parameter on the dynamical properties of the agent-environment coupling is substantiated by Figure 5. The bar-plot at the top of the figure shows the relative frequency of each of the nine attractors (labeled A to I corresponding to Figure 3) along the whole duration of the experiment (10 replications

Figure 4. Spectral analysis of the temporal activation of one of the frontal light sensors, during the occurrence of attractors type F (left) and H (right). The dominance of low frequency components and a more continuous spectrum at higher frequencies in the former, contrasted to the development of a highly periodic behavior (as shown by the well isolated spikes centered around the frequency of 0.45 Hz and its multiples) in the latter, allows a clear partition of their underlying dynamics. Detail of the interval [0 Hz, 2 Hz]; identical scales; analysis based on 1000 time steps after convergence to stable trajectories.

at 50 different levels of energy). In the lower part of the picture, the intensity of the pixels associated with each attractor details how the relative frequencies are distributed as a function of the current level of energy. Apparently, under the effect of the control parameter, the behavior drifts from attractors A and B, for the lowest levels of energy, to attractor I for energy levels of 0.8 and above. Attractor C prevails in the energy interval [0.1, 0.4], whereas the following interval is characterized by a sequence of couples or triplets of attractors appearing with similar probability for energy levels within the interval [0.5, 0.8] and slowly converging towards attractor H, stereotypical for high energies up to 0.8. The qualitative and quantitative nature of the observations, as resulting from the ten repetitions, received further validation by the results of additional repetitions testing different operative conditions (different levels of noise and random starting positions; details not shown here). Hence, the reported results, together with the following analysis, can be considered highly representative of the system's properties.

In accordance with Figure 5, the direct manipulation of the energy level, here used as a control parameter, effectively extracts a small subset of potential behavioral attractors, belonging to the pool of nine behavioral attractors embedded within the system by the evolutionary process, by shaping a specific landscape in the phase space. For example, applying a level of energy of 0.7, will activate either behavioral attractor H (70% of the trials) or, with a lower probability, F (20%) or D (10%). The actual attractor emerges on the basis of the actual starting position and of the integrated effects of noise, according to the general mathematical properties of dynamical systems. In this sense we can think of the energy level as an effective non-deterministic action selection mechanism, similar to but different from other suggested mechanisms (L. Cañamero, 2003; Prescott, Redgrave, & Gurney, 1999) in that this is emergent, dynamic and self-organized to a high degree.

Figure 5. Top graph: relative frequency of the nine observed behavioral attractors over all 500 replications. Bottom graph: the intensity of the pixels for each column (corresponding to attractors A-I) represents the relative frequency of the behavioral attractor for all ten replications at each level of energy. When the energy level is in the energy interval $[0.1, 0.4]$ we can observe a clear dominance of attractor ‘C’, expressed by the system in all the repetitions. A similar dominance is shown by attractor ‘I’ in the energy interval $[0.8, 1.0]$. An ordered sequence, involving two or more attractors expressed with similar probability at any energy level, characterizes the intermediate energy interval $[0.4, 0.8]$.

In the normal (evolved) conditions (see Figure 1), all the dynamical subsystems involved (internal energy level dynamic, neurocontroller, body and environment), are mutually engaged in a regulatory interaction, in order to achieve an effective and viable behavior. This mechanism, self-organized under evolutionary pressure, exploits the possibility to use the slower time scale of the energy level in order to integrate information over time (note that the neural controller is implemented as a purely reactive ANN). It should be emphasized that the energy level follows dynamics structurally dissimilar to the parallel interaction of the neurocontroller with its external environment (i.e. characterized by time scales that differ by orders of magnitude). As we can clearly see in Figure 5, when the agent suffers of an energy deficit (upper part of the lower graph) the pool of potential attractors is composed by “exploratory behaviors” (attractors A, B, C; cf. Figure 3), all looping between the two lights and engaging in external loops starting from one of them. Interestingly, attractor B, which expresses the possibility to engage in a limit cycle around the light source (although expressed with low probability), might be considered potentially maladaptive (inset in Figure 3, panel B). On the other hand, when the agent reaches an energy level around 0.8 or

slightly below (lower part of Figure 5) the potential attractors tend to engage in tight loops around the closest light source (attractors F, G, H). Attractors D and E integrate both characteristics, and could be considered transitional forms. Finally, when the energy level is above 0.8, the light source is interpreted as a mild repeller, as the agent tends to slowly step back from it. This mechanism reflects evolutionary constraints, in that the agent is not rewarded for stationary behaviors within the target area.

The spatial properties of the different families of attractors and the smooth transition from more exploratory to more local behaviors, are synthesized in more quantitative terms by the diagram in Figure 6. For each replication, we measured the minimum and maximum distance between the agent and both light sources. The diagram reports the variation of the average minimum and maximum distance (respectively, dark bars and solid line) as a function of the energy level. The average was based on the sequence of 10 repetitions at each energy level, as recorded during the clamp experiment (after all behavioral transients had settled). Actually, when the agent is engaged in local behaviors around a single light source, the maximum distance between agent and both light sources would be biased by the mutual distance between the two light sources. Therefore we decided to plot the minimum of the two maximum distances from each light source (minMax). Finally, open bars and dashed lines, in the same figure, report the respective standard deviations.

For energy in the interval $[0, 0.2]$, the average minimum-maximum distance is slightly higher than the mutual distance between the two lights. This correlates with the exploratory external loops in attractors A and B, partially mitigated by the overlapping expression of attractor C. As attractors C and D actually dominate the energy interval $[0.2, 0.6]$, the average minimum-maximum distance actually remains extremely stable, with a value close to the distance between the two lights (for the agent is moving from one to the next). The drift towards local attractors (F, G and H) in the interval $[0.7, 0.8]$ is represented by the smooth fall of this indicator. On the other hand, the average minimum distance regularly decreases in the interval $[0, 0.8]$. The final behavioral attractor of the sequence, H, is a quasi-circular movement whose radius tends to collapse to zero for growing values of energy level. As the energy level passes 0.8, a sudden, critical change affects the qualitative dynamics of the system. Both indicators in Figure 6, suddenly increase, and their similar values indicate the quasi-static nature of attractor I.

Following the structure of the neurocontroller (a simple feedforward network), the energy level directly modulates the motor activity. At the lowest energy levels, the two motor neurons receive highly negative activation, to the point of saturation (Figure 7). On the contrary, high levels of energy cause the inhibition of such vivid activation: in such conditions, the two motors tend to be activated at very low speeds, with a slight positive bias. Remarkably, according to a general mechanism that is well documented in the specialized literature (e.g. Nolfi, 2002), the evolutionary algorithm exploits contingent regularities: here in the form of the bodily symmetry of the simulated agent, that permits it to move in two opposite directions relying on two spatially distinct sets of sensors. Thus, the energy level modulates the balance between two opposite behavioral forces, phototaxis (at the lowest energy levels) and photophobia (at the highest ones), aiming for opposite directions and responsible for the generation of the overall dynamics of the system. Interestingly, the presence of such a push-pull mechanism is the basic framework for the

Figure 6. Variation of the minimum distance (dark bars) and of the minMax distance (see text; solid line) between the robotic agent and both light sources as a function of the energy level (averages over 10 replications for each level of energy). Open bars and dashed lines, drawn with the relative average as their baseline, represent the standard deviations.

emergence of deterministic chaos (Strogatz, 1994).

Validation under perturbed conditions.

In order to verify the dynamical nature of the behavioral attractors, we integrated our observations by running the energy clamp experiment in total absence of sensory noise. The expected collapse of the geometrical dimension of all behavioral attractors was indisputably contradicted by the experimental evidence. All attractors, except attractor I, re-exhibited a fractal geometry (for an example, see Figure 8, left)³. Attractor I turned into a fixed-point attractor.

To determine the robustness of the observed behavioral attractors (and possibly identify new ones) we tested the systematic injection of increasingly high levels of sensory noise. Interestingly, when sensory noise (uniformly distributed between 0% and 20% of the maximum possible activation of the sensor) was added to the current sensory input, neither the sequence of expression of the nine attractors, nor their relative frequency, were remarkably modified (details not shown here). Instead, all the behavioral attractors proved

Figure 7. Patterns of motor activation for different behavioral attractors at different levels of energy (plotted over 1000 time steps after extinction of the transient phase). The vertical and horizontal axes are associated with the activation level of the left and right motors. From left to right: energy level = 0.0 (attractor A); energy = 0.4 (attractor C); energy = 0.8 (attractor H). The pattern of activation tends to drift in a continuous fashion, as a function of the energy, from highly negative (at the lowest energy levels) towards a more bounded region with slightly positive average (for high energy).

persistent, and highly robust to perturbations. Nevertheless, interesting morphological peculiarities emerged: similar to attractor A (Figure 8, right), all the behavioral attractors are characterized by shaky and magnified trajectories, whereas the typical spectral content of random noise simply overlaps their usual spectra. Above this level of noise, a new critical transition seems to occur, as the sensors suddenly are no longer reliable enough to sustain coordinated behavior.

Discussion

On time scales and cognition

Theories of embodied and situated cognition stress the importance of the dynamical coupling between an agent and its environment. Furthermore, the body of any cognitive agent (biological or robotic) is much more than a mere physical substrate that allows mechanical interactions with the external world (Ziemke, Zlatev, & Frank, 2007; Chrisley & Ziemke, 2002; Clark, 1997). That means, the agent's body is not a cognitively neutral frame merely locating in space and time a handful of sensors and actuators for interaction. Rather, embodiment governs the specific modality of the agent-environment interactions, thus shaping the agent's cognitive potential. Accordingly, cognition does not happen in a vacuum, and any cognitive agent critically depends on and exploits the specificity of the properties of its environment. This effect has already been demonstrated in cognitive robotics (e.g. Nolfi & Floreano, 2000); brain, body and environment equally co-participate in the cognitive phenomena (cf. Chiel & Beer, 1997).

The work presented here is intended to contribute to this theoretical general framework by further exploring, and experimentally grounding, the basic principle of what Parisi called

Figure 8. Trajectories of the behavioral attractor A in the physical environment, respectively in total absence (left) and 20% (right) of sensory noise. In order to facilitate the comparison we maintained the same scale on both pictures, actually representing the whole extended arena.

“internal robotics” (Parisi, 2004): non-neural internal states, in interaction with all the other systems constitutive of the agent-environment coupling, are fundamental to the deployment of the full potential of embodied cognition. In our work, a simple non-neural internal state variable belonging to a simulated simple bio-regulatory process, the energy level, turned out to play a crucial role in regulating the behavioral coupling between the agent and its external environment.

We would like to emphasize that the coupling of systems with different time scales might be a key to the understanding of cognitive phenomena. The simple scenario used in the experiments discussed here, clearly demonstrates how, given the existing dynamics, high energy levels imply that a source of energy has been visited recently and therefore still must be in the neighborhood of the agent; on the other hand, low energy has the opposite meaning and thus results in exploratory behaviors. In this case, the representational role that traditional cognitive science tends to attribute to explicit memory, is totally offloaded on the global dynamic of the ‘brain-body-environment’ system. The slow dynamic of the energy level, coupled with the remaining elements of the overall system, implicitly triggers a mechanism functional to the integration of information over time, through a continuous, situation-dependent adaptation of the current robot’s behavioral repertoire. A closely related behavioral mechanism is reported in ecological contexts and modeled by Hills (Hills, 2006), where the dopaminergic modulation of area-restricted search strategy (ARS) is described with its evolutionary implications. We believe that, although admittedly very simple and limited, these results illustrate the fact that body-internal, bio-regulatory dynamics can play an important role in embodied cognitive phenomena (cf., e.g. Damasio, 1994; Lowe, Herrera, Morse, & Ziemke, 2008).

Within this context, the use of self-organized adaptation is not just a way to cope with the complex and generally unpredictable properties of the real environments. This is already technically problematic; in fact, even the mere possibility to describe declaratively viable and “intelligent” agent’s dynamics by a discrete set of rules, as typical of classical AI approaches,

is still a matter of intense scientific and philosophical debate. As others have argued before (Nolfi, 1998), from a cognitive or behavioral modeling perspective, self-organized adaptation also creates a value added by drastically reducing the impact of the designer's own ontology on constructing agent-environment dynamics. Therefore, self-organization tends to address and deploy in cognitive modeling the specificity of the cognitive potential of an agent-environment dynamical coupling, and its specific cognitive modalities.

Taken to the extreme, embodiment and situatedness imply that, once in place, self-organization tends to create classes of cognitive agents endowed with highly specific cognitive capacities. Consequently, an artificial cognitive agent can never be considered as a general model of cognition, but rather constitutes an indirect epistemic tool (a probe, you might say) by which we might investigate the basic principles behind cognitive phenomena. By observing the specific cognitive modalities in artificial models, we can infer basic information about the general organizational principles of biological embodied cognition, and the nature and relative roles of its basic elements. In that sense the emergence of self-organized dynamics in the context of interacting systems should be considered as the premise for the emergence of autonomous meaning.

In defense of the model

We would like to briefly point out three more things: Firstly, as mentioned above, the experimental scenario discussed here is of course simplistic, and it is intentionally so. In the choice of a body for our agent, the mechanism for its internal dynamics, its artificial nervous system and its environment, we systematically selected the simplest available option. Such a minimal setup allowed us to isolate and emphasize the phenomenon studied, and to extract/illustrate a general principle: the holistic nature of cognitive processes and the mutual contribution of internal and external dynamics to their emergence. Simplicity enabled us to carry out a detailed analysis of the ongoing phenomena. The use of more complex systems would, no doubt, give rise to more articulated dynamics; however, it would also be likely to blur the causal mechanisms to the extent of creating indistinguishable conceptual black boxes.

Secondly, our taxonomy of the observed behavioral attractors might be considered somewhat arbitrary. The classification of the attractors to some degree was based on their morphological characteristics, i.e. how their trajectories in the phase space 'appear' to us, the experimenters (whereas above, in fact, we argued for reducing the influence of the experimenter through self-organization). Furthermore, the spectral analysis of the attractors proved extremely useful in classifying otherwise morphologically similar attractors. It might be argued that some of the classes considered in Figure 3 could be merged into a single family, as their differences are merely due to the obvious time constraints of our observations. However, this does not affect our analysis in principle, since what is actually relevant here is the general structure of the transition between types/families of behavioral attractors that are qualitatively very different, respectively playing a major role at different energy levels, as demonstrated in Figure 6.

Finally, assuming that the dynamical systems approach to the study of embodied cognition might and should develop from its often somewhat metaphorical use of dynamical systems terminology, we tried to elucidate our analyses in mathematical terms, at least to some degree. In fact, the mathematical language is itself a metaphorical description of

the world (Lakoff & Núñez, 2000). Nevertheless, it represents a powerful and perhaps the only unambiguous conceptual tool our culture has produced for the accurate description of scientific theories and systematic exploration of their implementations as models. The same way logics offered a natural language to traditional symbol-systems-type AI, mathematics offers a natural language to the description of the dynamics and the emergence of embodied cognitive phenomena. That said, of course it remains to be seen exactly which mathematical tools and terminology turn out to be useful for the analysis of which types of cognitive and behavioral mechanisms. We believe this article has presented a couple of novel ways of viewing and analyzing minimal cognitive systems, and thus represents some progress towards the goal of developing a richer dynamical-systems toolbox for the analysis of embodied cognitive systems.

Conclusions and future work

We hope to have shown how a simple non-neural state, the energy level, can be used to modulate the dynamical properties of the interaction between a simple simulated robotic agent, controlled by a reactive feedforward artificial neural network, and its environment. This dynamical property, emergent and self-organized, turned out to play a crucial role in disambiguating perceptually identical light sources, on the basis of their capacity to influence the non-neural internal state. Nevertheless, the mechanism we analyzed should be considered as a general dynamical property and the specific instantiation merely prototypical for demonstration.

Crucially, the time scale of the energy level, very different from the time scale of the normal interactions between the agent and its environment, allows the system to integrate information over time. We argue that internal bio-regulatory dynamics in living systems can play a crucial role in supporting embodied cognitive processes. From this perspective, the non-neural internal state can also be considered an elementary, implicit mechanism of appraisal. Elsewhere (Herrera, Montebelli, & Ziemke, 2007; Herrera, Ziemke, & Moffat, 2006), we have argued in more detail that similar mechanisms represent the physiological basis for proto-emotions, in general accordance with somatic theories of emotions (Prinz, 2004; Damasio, 1994). From the self-organized coupling of internal dynamics, neurocontrollers, bodies and environments, all inter- and intra-operating on different time scales, might develop promising dynamics, from which genuine, autonomous, meaningful artificial cognition could emerge. Our major concern as cognitive system designers should be to incorporate in our models the critical level of complexity sufficient to trigger and sustain the process. Furthermore, the traditional roles that system theory attributes to controller and controlled, should be re-examined, as the nature of the mutual relationships among cognitive subsystems can only be contingent and circular. Even the cultural division that traditionally determines the boundary between inside and outside of the body, internal and external, should be critically analyzed in the perspective of cognitive modeling. Even if this distinction might seem quite straightforward in our current artificial agents, obviously developed according to our cultural biases, in biological agents such boundaries appear blurred and highly penetrable, for we are intrinsically extended, permeable, symbiotic biological machines (Turner, 2002; Clark, 2003).

The current study has focused on the role that non-neural internal (bodily) states might play in embodied cognitive phenomena. The internal dynamic we actually modeled,

as well as our crude manipulation of the control parameter, is of course a very rough abstraction from any dynamic normally found in natural bio-regulatory systems. Future work will be aimed at more detailed, multidimensional modeling of the non-neural internal (homeostatic) mechanisms and their role in the global dynamical interaction between cognitive agents and their environments. A refinement of the type of dynamical analysis presented here, aiming to identify the exact conditions for the occurrence of dynamical bifurcations and the dynamical sensitivity to perturbations, might also be relevant for a deeper comprehension of the general properties of embodied cognitive systems.

Acknowledgements

This work has been supported by a European Commission grant to the project “*Integrating Cognition, Emotion and Autonomy*” (ICEA, www.iceaproject.eu IST-027819) as part of the European *Cognitive Systems* initiative. The authors would like to thank Stefano Nolfi, developer of the Evorobot simulator used here, as well as Malin Aktius, Robert Lowe, Anthony Morse, Henrik Svensson and Domenico Parisi for their comments.

This article represents a substantial extension and revision (at experimental and theoretical level) of: Montebelli, Herrera & Ziemke, “An analysis of behavioral attractor dynamics”, presented at ECAL 2007.

Notes

¹Of course this internal mechanism, and the name itself, "energy level", represents an extremely crude simplification. What is important here is that this quantity is sensed by the agent, therefore influencing the activity of its artificial nervous system, and its dynamic was designed with a time scale that is slower by orders of magnitude than all other sensory-motor interactions.

²In dynamical system theory, a rigorous definition of the concept of attractor is somewhat elusive, particularly in most practical cases, when the problem at hand is not analytically well defined. Nevertheless, the use of the notion of attractor is here fully compatible with the list of the necessary and sufficient conditions suggested by some authors (cf., e.g. Strogatz, 1994). In particular all the attractors we classified are invariant, exhibit a basin of attraction and are minimal.

³We follow some authors (cf., e.g. Strogatz, 1994) in making a distinction between fractal and strange attractors. The former property, that merely focuses on geometrical aspects, does not imply the latter. A full characterization of the dynamical properties of the specified pool of attractors (for example their sensitive dependence on initial conditions, or an accurate description of the extension of their basins of attraction), or even the mere designation of the intrinsic reasons for the exhibited fractal trajectories, go beyond the scope of this article. In absence of a detailed analysis, we will simply avoid the trivialization of terms like strange attractor and chaos.

References

- Ashby, W. R. (1957). *Design for a brain: The origin of adaptive behavior*. London: Chapman & Hall.
- Avila-García, O., & Cañamero, L. (2004). Using hormonal feedback to modulate action selection in a competitive scenario. In S. Schaal, A. Ijspeert, S. Billard, & J. Vijayakumar (Eds.), *From animals to animats 8: Proceedings of the eighth international conference on simulation of adaptive behaviour* (p. 243-252). Cambridge, MA: MIT Press.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173-215.
- Cañamero, D. (1997). Modeling motivations and emotions as a basis for intelligent behavior. In W. L. Johnson (Ed.), *Proceedings of the first international conference on autonomous agents* (p. 148-155). New York, NY: ACM Press.
- Cañamero, L. (2003). Designing emotions for activity selection in autonomous agents. In R. Trappl, P. Petta, & S. Payr (Eds.), *Autonomous agents* (p. 115-148). Cambridge, MA: MIT Press.
- Chiel, H., & Beer, R. D. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20(12), 553-557.
- Chrisley, R., & Ziemke, T. (2002). Embodiment. In *Encyclopedia of cognitive science* (p. 1102-1108). London: McMillan.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (2003). *Natural-born cyborgs: Minds, technologies and the future of human intelligence*. New York, NY: Oxford University Press.
- Damasio, A. (1994). *Descartes'error: emotion, reason, and the human brain*. New York, NY: G.P. Putnam's Sons.
- Di Paolo, E. (2003). Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop. In K. Murase & T. Asakura (Eds.), *Dynamical systems approach to embodiment and sociality* (p. 19-42). Adelaide: Advanced Knowledge International.
- Gadano, S. C., & Hallam, J. (2001). Robot learning driven by emotions. *Adaptive Behavior*, 9(1), 42-64.
- Haykin, S., & Van Veen, B. (2003). *Signals and systems*. Upper Saddle River, NJ: Wiley.
- Herrera, C., Montebelli, A., & Ziemke, T. (2007). The role of internal states in the emergence of motivation and preference: a robotics approach. In A. Paiva, R. Prada, & R. W. Picard (Eds.), *Affective computing and intelligent interaction* (p. 739-740). Berlin: Springer.
- Herrera, C., Ziemke, T., & Moffat, D. (2006). Emotions as a bridge to the environment. In S. Nolfi et al. (Eds.), *From animals to animats 9: Proceedings of the ninth international conference on simulation of adaptive behaviour* (p. 3-16). Berlin: Springer Verlag.
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30, 3-41.
- 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.
- Ito, M., & Tani, J. (2004). On-line imitative interaction with a humanoid robot using a dynamic neural network model of a mirror system. *Adaptive Behavior*, 12(2), 93-115.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- Lakoff, G., & Núñez, R. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being*. New York, NY: Basic Books.
- 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.
- Maes, P. (1991). A bottom-up mechanism for behavior selection in an artificial creature. In J. A. Meyer & S. W. Wilson (Eds.), *From animals to animats: Proceedings of the first international conference on simulation of adaptive behavior* (p. 238-246). Cambridge, MA: MIT Press.
- Meyer, J. A. (1995). The animat approach to cognitive science. In H. L. Roitblat & J. A. Meyer (Eds.), *Comparative approaches to cognitive science* (p. 27-44). Cambridge, MA: MIT Press.
- Nicholls, J. G., Martin, A. R., Wallace, B. G., & Fuchs, P. A. (2001). *From neuron to brain*. Sunderland, MA: Sinauer Associates.
- Nolfi, S. (1997). Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior*, 5(3-4), 343-363.
- Nolfi, S. (1998). Evolutionary robotics: exploiting the full power of self-organization. *Connection Science*, 10(3-4), 167-183.
- Nolfi, S. (2002). Power and limits of reactive agents. *Neurocomputing*, 42(1-4), 119-145.
- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Parisi, D. (2004). Internal robotics. *Connection Science*, 16(4), 325-338.
- Pfeifer, R. (1999). Dynamics, morphology, and materials in the emergence of cognition. In S. Biundo, T. Frühwirth, & G. Palm (Eds.), *Advances in artificial intelligence: Proceedings of the 23rd annual german conference on artificial intelligence* (p. 27-44). Berlin: Springer.
- Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
- Prescott, T., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7(1), 99-127.
- Prinz, J. (2004). Embodied emotions. In R. C. Solomon (Ed.), *Thinking about feeling: Contemporary philosophers on the emotions* (p. 44-59). New York, NY: Oxford University Press.
- Strogatz, S. H. (1994). *Nonlinear dynamics and chaos*. Cambridge, MA: Westview Press.
- Tani, J. (2003). Learning to generate articulated behavior through the bottom-up and the top-down interaction processes. *Neural Networks*, 16, 11-23.
- Tani, J., & Ito, M. (2002). Self-organization of behavioral primitives as multiple attractor dynamics: A robot experiment. *IEEE Trans. on Systems, Man, and Cybernetics. Part A: Systems and Humans*, 33(4), 481-488.
- Turner, J. S. (2002). *The extended organism*. Cambridge, MA: Harvard University Press.
- Van Gelder, T. (2000). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615-628.
- Ziemke, T. (2000). On 'parts' and 'wholes' of adaptive behavior: Functional modularity and diachronic structure in recurrent neural robot controllers. In J. A. Meyer, A. Berthoz, D. Floreano, H. L. Roitblat, & S. W. Wilson (Eds.), *From animals to animats 6: Proceedings of the sixth international conference on simulation of adaptive behaviour* (p. 115-124). Cambridge, MA: MIT Press.
- Ziemke, T., Zlatev, J., & Frank, R. M. (Eds.). (2007). *Body, language and mind: Embodiment* (Vol. 1). Berlin/New York: Mouton de Gruyter.

Biographies

Alberto Montebelli

Alberto Montebelli is a Phd Student at the University of Skvde, Sweden, where he collaborates on the European funded ICEA project. He holds a “Laurea” degree in Electronic Engineering from the University of Bologna, Italy and a M.Sc. in Computer Science from the University of Skvde. His main research interests concern adaptive robotics, embodied and situated cognition, recurrent artificial neural networks and nonlinear dynamics.

Address: University of Skvde, School of Humanities & Informatics, PO Box 408, 54128 Skvde, Sweden. Email: alberto.montebelli@his.se

Carlos Herrera

Carlos Herrera is research associate in Cognitive Robotics at the Intelligent Systems Research Centre, University of Ulster. He holds a Licenciado degree in Mathematics (Universidad Complutense de Madrid), a Postgraduate Diploma in Philosophy (University of Glasgow) and a PhD (Glasgow Caledonian University) with the title The synthesis of emotion in Artificial Agents. Before moving to Ulster, he worked at the University of Skvde as a postdoc in the ICEA project 2006-2007. His main interests involve areas such as embodied, situated and dynamic approaches to cognition and adaptation, emotion, and the philosophy of robotics.

Address: University of Ulster, School of Computing and Intelligent Systems, Magee campus, Londonderry, BT48 7JL, UK. Email: c.herrera@ulster.ac.uk.

Tom Ziemke

Tom Ziemke is Professor of Cognitive Science at the University of Skvde, Sweden. His main research interests are theories and robotic models of situated and embodied cognition, in particular the various roles of the body in shaping cognitive processes. He is coordinator of a European project on bio-inspired cognitive/affective robotics (www.iceaproject.eu), editor of a book on embodiment recently published by Mouton de Gruyter (“Body, Language and Mind. Volume 1: Embodiment”), and associate editor of the journals “New Ideas in Psychology” and “Connection Science”.

Address: University of Skvde, School of Humanities & Informatics, PO Box 408, 54128 Skvde, Sweden. Email: tom.ziemke@his.se