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

This is an author produced version of a paper presented at 9th European Conference on Artificial Life, 10-14 September 2007 in Lisbon, Portugal. This paper has been peer-reviewed but may not include the final publisher proof-corrections or pagination.

Citation for the published paper:

Montebelli, Alberto; Herrera, Carlos & Ziemke, Tom

"An Analysis of Behavioral Attractor Dynamics"

Advances in Artificial Life, 2007 (Lecture Notes in Computer Science vol 4648) pp 213-222

ISBN: 978-3-540-74912-7

URL: http://dx.doi.org/10.1007/978-3-540-74913-4_22

Access to the published version may require subscription.

Published with permission from: Springer Verlag

An Analysis of Behavioral Attractor Dynamics

Alberto Montebelli, Carlos Herrera, Tom Ziemke

University of Skövde
School of Humanities and Informatics
SE-541 28 Skövde, Sweden
{alberto.montebelli, carlos.herrera, tom.ziemke}@his.se

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 notoriously difficult to analyze. In this paper 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 agent-environment system's behavioral attractors, such that the robot's behavioral repertoire is continually adapted to its current situation and energy level.

1 Introduction and Background

The dynamical systems (DS) approach to adaptive behavior and cognition (e.g. [1–3]) views agents and their environments as tightly coupled dynamical systems. Nervous systems, bodies, and their environments have complementary roles in producing a rich range of adaptive behaviors [4, 5]. Hence, the cognitive agent we consider is not a closed system 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-organising process. The implementation of artificial adaptation on the basis of principles of self-organisation is therefore alluring. The DS approach offers a range of tools available to the theorist that aims 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 [1, 6, 7]. Much recent robotics research has 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 plays an active role in control [7]. Furthermore, Parisi [8] 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 [9] 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" [10] - for the organism to function and stay alive. An adaptive agent is one that incorporates 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 [11, 12]. 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 [13–15]. 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. [16, 17]). 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 [17]. Similarly, in recent work on recurrent architectures with parametric bias values [18, 19], 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 paper, we further investigate the role of internal states in the modulation of behavior, exploring the hypothesis that non-neural internal states (metaphorically named "energy level") can modulate the functioning of a simple, recurrent neural network to produce different behaviors. Our perspective is radically different from related work on action selection since we do not consider behavior that is 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. Using a dynamical systems analysis, we distinguish several attractors, which can be understood as behavioral modes. The internal states used in the experiments documented here (energy values), are 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, therefore affecting behavior not through the selection of subsystems/modules, but through a modulation of the phase space that dynamically redraws the "behavioral repertoire" locally available to the agent.

2 Method

In the present paper 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 feedforward artificial neural network with no hidden layer, was artificially evolved by standard evolutionary algorithms within a square environment containing two identical light sources. The agent receives input from 8 infrared sensors, 8 light sensors and 1 energy sensor (tracking the current energy level), all directly connected, through the evolved weights, to the 2 motor neurons, respectively determining the activation of the left and right motor. The light sources are symmetrically located at the center of the arena, their reciprocal distance being comparable to the maximum distance range of the light sensors.

The robot was initially provided with an energy level of 1.0. Its energy was subject to a continuous decay at 0.008 per time step, down to zero. When the robot reached a target area surrounding one of the lights (highlighted, for the reader's benefit, by the circle on the left of Figure 1) its energy reservoir was instantaneously refilled, thus simulating an internal effect sensed by the agent. 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), under the condition that the agent were positioned outside the target area. On the other hand, the attainment of a null energy level would produce the opposite effect of integrating a negative constant (-0.02). In order to facilitate the evolutionary bootstrapping, the performance of the first 200 time steps of lifetime was not integrated into the fitness. The survival of the individual was only limited by the number of time steps during its lifetime and immediately obliterated by eventual crashes against the walls, which also determined a severe punishment (-100).

Fig. 1. Evolved behavior. The simulated agent approaches the light (represented by a dot) on the right side of the environment. The lack of any effect on its energy level sensor results in the following approach of the next light. As its energy reservoir is instantaneously refilled to the maximum level, the agent engages in a stable behavioral attractor in the proximity of the rewarding light source. The signals labeled LS1-8 represent the activation of the light sensors in time; EnS displays the energy level.

3 Results

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, quickly migrates towards the other light source that supplies it with a whole load of energy. Such an offer produces the engagements of the agent in a stable behavioral attractor in the neighborhood of the rewarding light source. The development of a skillful controller capable of light disambiguation constitutes, rather than the goal, the actual starting point for our work, which will emphasize the general holistic nature of the relation between agents and their environments, the complex dynamics emerging from such interactions and the peculiar 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 agent, its control system, its environment/arena and the set of bio-regulatory mechanisms (synthesized, as a first approximation, in the agent's energy level). In the following analysis the energy level is used as control parameter of the overall system. The systematic manipulation of the energy level allows us to visit a wealth of behavioral attractors embedded in the system during its evolution (where the adaptation is limited to its neural controller).

3.1 Energy Binary-Switch

We ran some simulations where the same agent as 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 to low. In this case none of the two light sources would provide the agent with any energy recharge.

As shown in Figure 2, during the first 300 time steps the agent drifts at low speed away from both light sources, moving from its initial position, labeled 'a', towards position 'b'. As the energy level switches from 1.0 down to 0.0 (emphasized by an arrow in the signal plots), the speed is abruptly and remarkably increased (as clearly shown by the much longer distance run in only 200 time steps compared to the distance from 'a' to 'b') and the agent engages in a fast run, alternating visits to both of the light sources, in its vane quest for energy.

3.2 Energy Clamp

We further developed the analysis by refining the quantization of the possible energy levels. The previous controller was fed with discrete, steady levels of energy, ranging from 0.0 to 1.0 with constant increments of 0.02, and the agent left free to interact for a time sufficient to fully develop its ongoing dynamic. We

Fig. 2. Energy binary switch experiment. During the first 300 time steps (energy level 1.0), the same agent as illustrated in Figure 1 slowly moves from the position labeled 'a' towards 'b'. When the energy level switches to zero (arrow in the plot of the sensory activations) its speed suddenly increases, as the agent alternately visits the two lights.

developed a taxonomy constituted of nine observed behavioral attractors, based on their general morphology and spectral characteristics (Figure 3), in order to determine the effect of the control parameter on the dynamical properties of the agent-environment coupling (Figure 4).

Fig. 3. Iconic representation of the spatial trajectories for the nine behavioral attractors in our taxonomy. The scale has been systematically varied for the reader's benefit, in order to facilitate a deeper qualitative grasp of the attractor's morphology (where possible, light sources are indicated by a star).

The bar plot at the top of Figure 4 shows the relative frequency of each of the nine attractors along the whole duration of the experiment. The lower part details how such relative frequency is 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 energies of 0.8 and above. The energy interval $[0.1, 0.4]$ is dominated by attractor C, whereas the following interval is characterized by a sequence of couples or triplets of attractors appearing with similar probability for energy levels within $[0.5, 0.8]$ and slowly converging towards the stereotypical attractor for high energies.

Fig. 4. Top graph: relative frequency of the nine observed behavioral attractors over all replications. Bottom graph: the intensity of the pixels in each column (respectively associated with attractors A-I) represents the relative frequency of the behavioral attractor for each level of energy.

4 Analysis and Discussion

The tenets of embodied and situated cognition stress the importance of the physical dynamical coupling between an agent and its environment. This effect has already been shown in cognitive robotics (e.g. [20]); brain, body and environment equally co-participate in the cognitive phenomena. With the present work we contribute to such general framework by adding, and grounding in experimental terms, the basic principle of internal robotics: non-neural internal states are fundamental to the deployment of the full potential of cognitive processes. In our work a simple non-neural internal state variable belonging to the internal bio-regulatory process, the energy level, proves effective at regulating the behavioral coupling between the agent and its external environment.

When the energy level is used as a control parameter (see previous section) it effectively evokes a subset of potential behavioral attractors, belonging to the pool of behavioral attractors embedded within the system during the evolutionary process, by shaping a specific landscape in the state space. According to

the general mathematical properties of dynamical systems, the actual attractor then emerges on the basis of the actual starting position and the current effects of noise. In this sense the energy level can be effectively used to implement a sort of non-deterministic action selection mechanism, similar to but different from other suggested mechanisms [14, 21] in that this is emergent, dynamic and self-organized to a high degree.

In the normal (evolved) conditions (see Figure 1), all the dynamical subsystems (internal energy level dynamic, neurocontroller, body and environment) are engaged in a mutual regulatory interaction, in order to achieve an effective and viable behavior, self-organized under evolutionary pressure. During the normal dynamic interaction of the bio-regulatory and agent-environment subsystems, the mutual regulation exploits the possibility to use the energy level in order to integrate information over time (remember 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 different time scales). As we can clearly see in Figure 4, when the agent is in a condition of energy deficit (upper part of the lower graph) the pool of potential attractors is composed by "exploratory behaviors" (attractors A, B, C), both looping between the two lights and engaging in external loops starting from one of them. Interestingly, attractor B, for it embeds the possibility to engage in a limit cycle around the light source (although expressed with low probability), might be considered potentially maladaptive. On the other hand, when the agent reaches an energy level around 0.8 or slightly below (lower part of the figure) the potential attractors tend to engage in tight loops around the closest light source (attractors F, G, H). Attractors D and E embed both characteristics, and might be considered transitional forms. Therefore, given the dynamics in place, high energy level implies that a source of energy was recently visited and is still in the neighborhood of the agent, whilst low energy has an opposite value and is compatible with exploratory behaviors. In this case, the role that traditional cognitive science tends to attribute to explicit memory, is totally offloaded on the global dynamic of the system. Accordingly, internal bio-regulatory dynamics might be interpreted as supportive of cognitive phenomena.

As mentioned above, we intentionally carried out our experiment in a simple scenario. Such a minimal setup was functional to isolate and emphasize the phenomenon we intended to study: the holistic nature of cognitive processes and the mutual relationship of internal and external dynamics in their emergence. A more complex system might indeed give rise to more complex dynamics; nevertheless it would blur its causal mechanisms to the extent of creating an indistinguishable conceptual black box. On the contrary, we are allowed to develop a detailed analysis of the ongoing phenomena. For example, it is quite interesting to observe how the pattern of motor activation is modulated by the energy level. At the lowest energy levels, the two motor neurons receive high negative activation, to the point of saturation (Figure 5). On the contrary, high levels of energy cause the inhibition of such vivid activation: in such conditions, the two motors

Fig. 5. Patterns of motor activation for different behavioral attractors at different levels of energy. 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 as a function of the energy, from highly negative (at the lowest energy levels) towards a bounded region of slightly positive values (for high energy).

tend to be activated at very low speeds, with a slight positive bias. Therefore, according to a mechanism that is well documented in the specialized literature, the evolutionary algorithm exploits contingent regularities, here in the form of bodily symmetry of the simulated agent that permits it to move in two opposite directions, at the same time relying on two rather well 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. Incidentally, such a push-pull mechanism is the basic framework for the emergence of deterministic chaos [22].

A final clarification regards our taxonomy of the behavioral attractors, which might be considered somewhat arbitrary. The classification of the attractors was based on their morphological characteristics. Furthermore, their spectral analysis proved extremely useful in order to classify otherwise morphologically similar attractors. For example, attractors F and H, sharing the property of being bound to a limited region in the neighborhood of one of the lights, can be easily disambiguated on the basis of the spectral information carried by the activation of the light sensors (Figure 6). It might be argued that some of the classes considered in Figure 3 might be merged into a single family, as their differences are merely due to the obvious time constraints of our observations. Nevertheless, such a remark would not deeply affect our analysis in principle, since what is actually relevant to our purpose is the general structure of the transition between families of attractors that are qualitatively very different, respectively playing a major role at lower or higher energies. A quantitative analysis of the behavioral attractors might provide a detailed analysis of the occurring dynamical bifurcations and of the sensitivity of the system to local conditions (initial state, noise, etc.), beyond the intentions of the current work.

5 Conclusions and Future Work

With the present study, we show how a simple non-neural state belonging to internal bio-regulatory processes, the energy level, can be used to modulate the

Fig. 6. Spectral analysis of the temporal activation of one of the light sensors, during the occurrence of attractors type F (left) and H (right). The dominance of low frequency components and a more continuous spectrum in the former, contrasted to the development of a highly periodic behavior in the latter (as shown by the isolated spikes centered around the frequency of 0.5 Hz and its multiples), allows a clear partition of their underlying dynamics. Detail of the interval [0 Hz, 2 Hz]; identical scales.

dynamical properties of the interaction between a simple simulated robotic agent, controlled by a feedforward artificial neural network and its environment. This dynamical property, emergent and self-organized, is here functional in the task of disambiguating identical light sources, on the basis of their capacity to influence the non-neural internal state. Nevertheless, the mechanism we describe should be considered as a general dynamical property and the specific problem merely prototypical for demonstration.

Crucially, the time scale of the energy level, peculiarly different from the time scale of the normal interactions between the agent and its environment, allows the neural controller of the agent to integrate information over time. We argue that internal bio-regulatory dynamics might play a role in supporting cognitive processes. Under this perspective, the non-neural internal state can also be considered as the implementation of an elementary, implicit mechanism of appraisal. In a complementary work [23] we argue that similar mechanisms represent the physiological basis for proto-emotions, in general accordance with somatic theories of emotions [24, 25].

The current study has focused on the role that non-neural internal (bodily) states might play in support of cognitive phenomena. The internal dynamic we actually modeled, as well as our crude manipulation of the control parameter, is indeed a very rough simplification with respect to the normal dynamic of the bio-regulatory systems observed in nature. In future works we will try to understand how a more detailed modeling of the non-neural internal mechanisms (i.e. modeling of internal homeostasis) might further enhance the potential of the global dynamical interaction between cognitive agents and their environments. The refinement of our dynamical analysis, so 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 system.

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.

References

1. Clark, A.: Being There. MIT Press (1997)
2. van Gelder, T.: The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* **21** (1998) 615–628
3. Kelso, J.A.S.: Dynamic Patterns. MIT Press (1995)
4. Beer, R.D.: A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* **72** (1995) 173–215
5. Chiel, H., Beer, R.: The brain has a body. *Trends in Neurosciences* **20**(12) (1997) 553–557
6. Pfeifer, R., Scheier, C.: Understanding Intelligence. MIT Press (1999)
7. Pfeifer, R.: Dynamics, morphology, and materials in the emergence of cognition. In: *Advances in Artificial Intelligence: Proceedings of the 23rd Annual German Conference on Artificial Intelligence*, Springer (1999) 27–44
8. Parisi, D.: Internal robotics. *Connection Science* **16**(4) (2004) 325–338
9. Ashby, W.R.: Design for a Brain. Chapman I& Hall (1960)
10. Meyer, J.A.: The animat approach to cognitive science. In Roitblat, H.L., et al., eds.: *Comparative Approaches to Cognitive Science*. MIT Press (1995) 27–44
11. Avila-García, O., Cañamero, L.: Using hormonal feedback to modulate action selection in a competitive scenario. In Schaal, S., et al., eds.: *From animals to animats 8*, MIT Press (2004) 243–252
12. Gadanho, S.C., Hallam, J.: Robot learning driven by emotions. *Adaptive Behavior* **9**(1) (2001) 42–64
13. Cañamero, D.: Modeling motivations and emotions as a basis for intelligent behavior. In: *Proceedings of the first international conference on Autonomous agents*, ACM Press (1997) 148–155
14. Cañamero, L.: Designing emotions for activity selection in autonomous agents. In Trappl, R., et al., eds.: *Autonomous Agents*. MIT Press (2003) 115–148
15. Maes, P.: A bottom-up mechanism for behavior selection in an artificial creature. In Meyer, J.A., et al., eds.: *From animals to animats*, MIT Press (1991) 238–246
16. Nolfi, S.: Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior* **5**(3-4) (1997) 343–363
17. Ziemke, T.: On 'parts' and 'wholes' of adaptive behavior. In: *From animals to animats 6*. (2000)
18. Ito, M., Tani, J.: On-line imitative interaction with a humanoid robot using a dynamic neural network model of a mirror system. *Adaptive Behavior* **12**(2) (2004) 93–115
19. Tani, J.: Learning to generate articulated behavior through the bottom-up and the top-down interaction processes. *Neural Networks* **16** (2003) 11–23
20. Nolfi, S., Floreano, D.: *Evolutionary Robotics*. MIT Press (2000)
21. Prescott, T., Redgrave, P., Gurney, K.: Layered control architectures in robots and vertebrates. *Adaptive Behavior* **7**(1) (1999) 99–127
22. Strogatz, S.H.: *Nonlinear Dynamics and Chaos*. Westview Press (1994)
23. Herrera, C., Ziemke, T., Moffat, D.: Emotions as a bridge to the environment. In Nolfi, S., et al., eds.: *From animals to animats 9*, Springer Verlag (2006) 3–16
24. Damasio, A.: *Descartes'error*. G.P. Putnam's Sons (1994)
25. Prinz, J.: Embodied emotions. In Solomon, R.C., ed.: *Thinking about Feeling*. Oxford University Press (2004) 44–59