Chapter 7
On the Role of Embodiment for Self-Organizing Robots: Behavior As Broken Symmetry

Ralf Der

7.1 Introduction

Embodiment and SO form two cornerstones of both modern robotics and the understanding of human and animal intelligence. In particular, the role of the embodiment for the behavior of both artificial and natural beings has become of much and increasing interest in recent times. In robotics, there are essentially two attitudes towards the physical embodiment. On the one hand, with rule based systems and/or systems intended to execute a given motion plan, embodiment is more or less considered as a (nasty) problem opposing the execution of the plan. On the other hand, it is well believed and verified by many examples that living beings are taking much advantage from the physico-mechanical properties of their bodies in order to create natural motion patterns. In robotics, this can be of immediate benefit for more robust or energy effective control, The importance of the embodiment for the behavior generation and intelligence in general has been advocated with great impact on the scientific community mainly by the lab of Rolf Pfeifer, see (Pfeifer and Bongard 2006; Pfeifer et al. 2007; Pfeifer and Gomez 1999) for excellent surveys. In recent years this has been further established under the connotation of morphological computation (Pfeifer and Gómez 2009; Hauser et al. 2012; Pfeifer 2012; Hauser et al. 2011). Seminal contributions are also by the lab of Helge Ritter (Ritter et al. 2009; Grossekathofer et al. 2011; Behnisch et al. 2011; Elbrechter et al. 2011; Steffen et al. 2010; Maycock et al. 2010) and others.

Self-Organization (SO) may provide an essential progress in the realization of embodied control. Viewing a robot in its environment as a complex dynamical system, SO can help to let highly coordinated and low dimensional modes emerge in the coupled system of brain, body and environment. In this way, instead of being programmed for solving a specific task, the robot may find out by itself what its bodily

Ralf Der
Max Planck Institute for Mathematics in the Sciences, Inselstraße 22, D-04103, Leipzig, Germany
e-mail: ralfder@mis.mpg.de

M. Prokopenko (ed.), *Guided Self-Organization: Inception,* Emergence, Complexity and Computation 9,
DOI: 10.1007/978-3-642-53734-9_7,

affordances are, focusing only in a second step on the exploitation of the emerging motion patterns—by guiding the SO process into the directions of potential benefits.

This paper studies the relation of embodiment to self-organization in autonomous robots. This project faces essentially two challenges. One is how to organize a robotic system in such a way that it starts to self-organize. There are several approaches into that direction based on formulating objective functions (OFs) for SO. In recent years, several such OFs have been proposed, ranging from the maximization of predictive information (Ay et al. 2008, 2012; Martius et al. 2013) or empowerment (Klyubin et al. 2005, 2007; Anthony et al. 2009; Jung et al. 2012), to the minimization of free energy (Friston 2010; Friston et al. 2012; Friston and Stephan 2007; Friston 2012) or the so called time-loop error in the homeokinesis approach (Der 2001; Der and Liebscher 2002; Der and Martius 2012), see also (Prokopenko 2008, 2009; Prokopenko et al. 2009) for more details on how to organize SO. Given an objective function, the optimization process can be translated into a learning rule that is driving the SO process. This paper introduces a new learning rule together with some case studies demonstrating its usefulness. So, there are several pretty satisfying approaches in the SO paradigm.

The second challenge actually is more serious and might be the reason why in the community SO is seen more as a wishful thinking than a systematic approach to autonomous robot development. With genuine SO one must be careful not to plug in (by biasing the system) what one actually wants to get out. The approaches mentioned above all seem to fit that criterion. However, and this seems to be the argument, if nothing is specified from outside, will SO simply make the robot an arbitrary subject that is completely unpredictable in its behaviors and thus rather a thread than a hope. The aim of this paper is to show that this attitude is wrong. Instead, we will develop an understanding of what happens if the system is self-organizing, what the role of the embodiment is and how we can find clues for predicting and shaping the behavior patterns emerging in a genuine SO scenario.

In order to prepare the ground for the role of the embodiment in SO, we give some examples of systems unfolding complex motion patterns with minimalistic control. We start from the famous example of the Braitenberg vehicles (BVs) which we consider as an early case study of machines which develop complicated behaviors on the basis of an extremely simple construction. By these examples, we want to make the reader aware of the phenomenon of spontaneous symmetry breaking that in our opinion is instrumental for understanding how SO can be effective in robotic systems. We think that the robotic community so far has overlooked the importance and substance of that phenomenon. Therefore, this paper will also follow a pedagogical purpose.

The paper is organized as follows. After introducing the original idea of Valentin Braitenberg and variations thereof in Sect. 7.2, we introduce in Sect. 7.3 a much more complex body, our HUMANOID, that develops a complex behavior mode under a minimalistic control. This gives us the opportunity to introduce the concept of fundamental modes, dynamical patterns that are specific for a particular body under homogeneous energy feeding condition. This is also a first example of spontaneous symmetry breaking (SSB). Sect. 7.4 introduces the new learning rules for driving

SO based on a recent result from maximizing predictive information (Martius et al. 2013), followed by a short aside on the principle of homeokinesis in Sect. 7.5 containing a comparison between the two methods. Sect. 7.6 is devoted to a deeper discussion of fundamental modes and the mechanism of SSB. After discussing these phenomena with the example of the autistic vehicles, Sect. 7.7 and Sect. 7.8 study these questions with the HUMANOID and the HEXAPOD robot.

7.2 Vehicles

Let us start with a few ideas on what we want to understand by embodiment and its synergy with the "brain". Let us go back to a very instructive example given by Braitenberg (Braitenberg 1984) in order to better understand the synergy concept.

7.2.1 Braitenbergs Idea

His idea was to use very simple machines in order to demonstrate the emergence of complex behaviors determined by the very physical construction of those machines. In the most simple case, the machine is a two-wheeled robot with two photo cells mounted to the left an right front side of the trunk. The photocells are wired to the motors either diagonally so that the vehicle is approaching the light source and bangs into it, or directly so that light sources are acting as a repeller.

Besides the phenomenon of emerging functionality observed in those machines, there is another fact we want to draw attention to. Both in the basic and also in the more elaborate examples, it seems more or less artificial to make a distinction, as is done in most robotics approaches, between the body and the "brain". Instead, we see a physical system—a mechanical part consisting of the body and the wheels combined with an electrical part consisting of the photo cells, the motors and the wires. With this setting it seems more natural to see the machine as a whole, instead of subdividing it into a body and a controller acting as a (kind of) brain. Of course, which attitude to take is always a matter of taste, preferences shifting more toward a clear partitioning with increasing complexity of the "brain". This paper, keeping control as simple as possible, is more in favor of the holistic attitude: considering the machine—body, sensors, and wiring—as one physical system. In the Braitenberg case, this system may function even completely autonomously since it receives all the energy it needs from its sensors, the photocells.

7.2.2 Autistic Vehicles

The properties of the most simple vehicles are self-explaining given that the behavior is driven by the energy provided by the photo cells together with the specific wiring. If translated to humans or higher animals, this corresponds to a system driven by vision alone. However, motion control is largely depending also on the signals they are receiving from proprioception, like the joint angles and the muscle

tensions. In order to mimic the role of those sensors, we are going now to restructure the BVs by including proprioceptive sensors. In particular, we introduce wheel counters, measuring the rotational velocity of each of the wheels. Wiring can be done in the most simple case by connecting the right (left) wheel sensor to the right (left) motor and vice versa. However, we do now need active wires which feed energy into the system (there is no perpetual motion machine of the first kind).

Such an active, nonlinear wire can be realized by a simple neuron, that translates the input (wheel velocity) into an output— the target wheel velocity that is subsequently realized by the motor using electrical energy supplied from outside. In detail, we use a neuron as

$$y = \tanh(cx + h) \tag{7.1}$$

where c is the coupling strength and h a threshold to be put equal to zero for the moment. This setting has been studied before in different contexts. The behavior of the wheel can be seen if we consider the full sensorimotor loop (SML). What we need for that is the connection between the target motor values y and the corresponding wheel velocity x as measured by the wheel counter. Assuming a linear relation between the two, the SML can be modeled by the following dynamical system.

$$x' = ay + \xi = a\tanh(cx + h) + \xi \tag{7.2}$$

where x' is the new vector of sensor values, ξ is to contain all the deviations from the linear law, and a is a hardware constant, with $a = 1$ if the wheel counter is scaled appropriately. In many applications, ξ can be treated as pure noise.

The behavior of the vehicle can be analyzed best by asking for the fixed points (FPs) of the dynamical system which are obtained from $x = ay$ or (with $a = 1$)

$$x = \tanh(cx) \tag{7.3}$$

There are a number of different regimes. Let us consider first the case $0 < c < 1$ with the single FP $x = 0$ corresponding to wheel velocity equal to zero. With noise, each wheel velocity is fluctuating around zero so that the vehicle executes a random walk. With $c > 1$ the system is bistable. If $c = 1 + \delta$, with $\delta > 0$ small, the FP is given approximately as $x = \pm\sqrt{\delta}$, see (Der and Martius 2012) for details. In this case, the system is in one of the FPs but can be switched by the noise. So, the wheel is rotating for some time into one direction with occasional changes of direction. This effect generates an irregular motion of the vehicle due to the noise. In particular, upon colliding with an obstacle, the velocity may switch so that the vehicle is kind of reflected by the obstacles. This may be seen as a basic survival strategy emerging from this minimalistic control. In this sense, the autistic vehicle still does show certain reactions to the environment. However, these are restricted to noise or physical encounters and the perturbations generated thereby.

With $c < 0$, the output of the neuron (motor command) is always opposite in sign to its input so that the motor always gets contradictory commands which might eventually lead to the destruction of the motor. Therefore, this regime was called the regime of self-destruction. However, we will see in Sect. 7.3 below that, based on

the specific physical embodiment, this regime may also deploy highly interesting modes.

Altogether there are not much interesting properties of that autistic vehicle. A first interesting result is observed if several such two-wheeled vehicles are coupled passively together. In that case, the physical cross-talk—the coupling forces produced by different wheel velocities of neighboring vehicles—may switch the wheel velocities so that eventually a collective motion of the system emerges. This has been studied extensively earlier (Der et al. 2008; Zahedi et al. 2010; Der and Martius 2012) so that we do not go into details here. Instead we consider other more complex systems where a collective motion is emerging from decentralized control in a surprising way, see Sect. 7.3 below.

7.2.3 Symmetries

Symmetries and their breaking play a central role in this approach. Let us start here with the standard scenario of inducing the breaking of a given symmetry by manually driving certain parameters of the system over a bifurcation point. The self-induced breaking that is of actual interest for this paper will be discussed in Sect. 7.6 below. By way of example, let us consider the autistic vehicle introduced above with a single wheel running on a rail. Assuming there is a perfect forward-backward symmetry of the morphology, the system is invariant against inversion of the x-axis, i. e. it obeys the symmetry $S : x \to -x$.

Given that the controller is invariant against S and assuming that the noise also is, we find that there is no clue for the controlled system to prefer a specific direction in space[1]. In the regime with only a single FP (0<c<1), this means that the trajectory is fluctuating around the FP $x = 0$. In the sense of a classical bifurcation analysis, we can now increase the coupling strength c. When crossing the bifurcation point, the system is entering its bistable regime with two stable attractors corresponding to the wheel rotating forward or backward. Which one is chosen is determined by the noise so that we have here a trivial case of noise-induced symmetry breaking.

7.3 The Braitenberg Man—Fundamental Modes

Let us now extend Braitenberg's idea to a more complicated machine—our HUMANOID with a specific wiring.

7.3.1 The HUMANOID

We consider a humanoid robot with 17 active degrees of freedom. Each joint is driven by a simulated servo motor, the motor values $y \in \mathbb{R}^{17}$ sent by the controller

[1] With noise the symmetry is to be understood in the stochastic sense, meaning that each trajectory finds its counterpart—a trajectory that is realized with the same probability—by applying the symmetry operation.

being the target angles of the joints and the sensor values $x \in \mathbb{R}^{17}$ are the true, i. e. observed angles. This is the only knowledge the robot has about its physical state. The physics of the robot is simulated realistically in the LPZROBOTS simulator. Let us introduce a specific wiring in analogy to the autistic vehicle: each joint i is controlled by a single neuron generating the control $y_i = \tanh(cx_i)$ for motor i, with c a constant coupling strength, the same for all joints. In the terminology of Braitenberg the neuron can be considered as an active, nonlinear wire. By this wiring, like in the case of the coupled vehicles mentioned above, we have a system with fully decentralized control (called split control in (Zahedi et al. 2010)).

7.3.2 *A Fundamental Mode of the* HUMANOID

The interesting behavior emerges if we give the robot much freedom to develop its specific, body-inspired motions. In the experiment, we suspend the HUMANOID like a bungee jumper. After starting it in a position slightly above the equilibrium point of the robot-spring system, the robot is seen to fall down a short distance and then oscillates vertically for a while due to the elastic forces of the spring, without developing internal motions of its own. Its further fate depends essentially on the coupling strength c. We do not observe much of interest for positive coupling strengths. Instead, there is a critical negative value c_{crit} so that, with $c < c_{crit}$, the robot is seen in the experiment—after the oscillatory phase—pretty soon to go into a regular, periodic motion[2]. This is caused by the self-amplification of initial, very small perturbations, see the video S1 (in the supplementary material on `http://robot.informatik.uni-leipzig.de/research/supplementary/GSO2012/`). The interesting and on first sight surprising observation is that this motion is very similar to a running behavior. In particular, despite of the completely decentralized control, arms, legs and hip motions are seen to be synchronized with fixed phase relations like with a running human.

Why is that? How do the individual controller know of each other? In order to discuss this point we must remember that the coupling constant is strongly negative so that actually we are in the regime of self-destruction. However, in the physical system, motor forces are confined to a low range so that—given the inertia due to the masses of the segments—the joint angles can change only slowly. Thus, as long as a joint angle is largely negative, a strong positive motor command is generated that engenders a motion towards the region of positive joint angles. Overshooting (by both inertia and the time averaging of the sensor values) the region of small joint angles, the scenario repeats with an impact towards negative joint angles. Altogether this produces an oscillatory motion of the joint with frequency and amplitude depending on the motor forces and masses of the moved segments in a complicated way.

This effect explains the latent oscillatory mode of each individual joint. But why the (running like) coordination between the segments? Well, the motions are seen to be quite rapid so that there are heavy intra-robot physical forces: by the inertia

[2] We have to smooth the sensor values in a short time horizon to obtain that result.

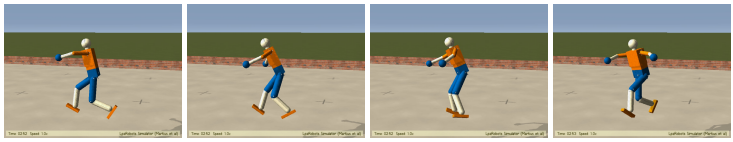

Fig. 7.1 One half-period in the running motion pattern of the Braitenberg man. In the simulation we use a time average over 15 steps of the sensor values.

effects, any motor activity changing a joint angle will generate a heavy impact on its neighbors. So, like with the chain of wheeled robots, there is a strong physical cross-talk that is the communication basis of the intra-robot coordination phenomenon. Spontaneous symmetry breaking (SSB) comes into play here. By its geometry, the specific wiring, and the identity of all the coupling constants, the system is highly symmetric, any emerging motion pattern corresponding to a specific breaking of that symmetry. Due to the overcritical negative feed-back strength given by $c < c_{crit}$, there is a strong tendency of self-amplification of the physical cross talk. Given the high degree of symmetry at the start, this self-amplification induces a symmetry breaking with a specific motion pattern emerging. Choosing the perturbations by hand, many different motion patterns are possible but in the experiments, the running motion pattern (RMP) is dominating, i. e. it seems to be the most stable one—once the pattern is established one may perturb it by external forces but, in most cases, the pattern is recovered after a short time.

The above mechanism also explains why we want to call the emerging dynamics a fundamental mode of the system. We borrowed that term from physics where fundamental modes (often called normal modes) are common phenomena. For instance, in coupled systems of linear or nonlinear oscillators, low frequency modes are emerging where all oscillators are synchronized with a fixed phase relation. In neuroscience such modes are also known for networks of integrate-and-fire neurons, e. g. Synchrony and fixed phase relations are also featuring in the RMP—the individual constituents of the body are moving in a highly coordinated manner. Moreover, the mode may be called fundamental since it is a dynamical pattern that is specific for this particular body under the homogeneous energy feeding condition: the intake of energy is regulated for each motor in the same way by an individual feed-back loop with the same coupling constant. Yet, despite this structural homogeneity we get a highly structured dynamics.

The existence of the fundamental mode heavily relies on the specific conditions, in particular the freedom of movement as realized in the case of the bungee jumper. When lying on the ground, the robot will not develop any interesting activity. With most values of $|c| < c_{crit}$, the robot will reach a fixed point behavior, c_{crit} corresponding to the overall critical feed-back strength. With super-critical values, the motions are confined by the contact with the ground. These restrictions are overcome with a convenient method for self-organization of the behavior as we will develop it in the following.

7.4 Unsupervised Learning for Self-Organization

So far, we have organized the system for generating specific modes by hand, guided by an intuitive understanding of the physics and the physical cross-talk mechanisms. Let us now start with developing a scheme for the self-organization of the system, hoping that the system itself discovers such interesting phenomena like the existence of fundamental modes.

In recent work, the so called predictive information (PI) was introduced as a general objective function for SO (Ay et al. 2008, 2012; Zahedi et al. 2010; Martius et al. 2013). By maximizing the PI, a general learning rule for the synaptic strengths of a neural controller network was derived on the basis of such an Infomax principle. Different from Infomax principles derived so far, that method solves the task of relating the general principle, formulated at the level of behavior, to the internal world of the robot, in to its brain so to say.

The controller of the robot is given in our case by an artificial neural network (ANN) transforming sensor values $x \in \mathbb{R}^n$ into motor values $y \in \mathbb{R}^m$ like

$$y = K(x, ..., w) \tag{7.4}$$

where w are the parameters (synaptic strengths) and ... means some other, internal variables which we will drop for the moment. The translation between the external and the internal world can be done if there is a forward model predicting future sensor values on the basis of the current sensor and motor values. In particular we need only the relation between two time steps so that

$$x_{t+1} = \phi(x_t, y_t) + \xi_{t+1}$$

where ξ is the prediction error and the parametrized function $\phi : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is the predictor.

7.4.1 Learning Rules for Self-model and Control

We consider here only the result for a specific controller—a one layer feed-forward neural network realized as

$$K(x) = g(Cx + h) \tag{7.5}$$

so that the set of parameters $w = (C, h)$ is given by the synaptic values (matrix C) and the vector h of threshold values for the neurons. In the concrete applications done in this paper, we specifically use $g_i(z) = \tanh(z_i)$ (so that $g : \mathbb{R}^m \to \mathbb{R}^m$ is to be understood as a vector function).

Moreover, the forward model ϕ is given by a layer of linear neurons, so that

$$\phi(x, y) = Ay + Sx + b\,. \tag{7.6}$$

The matrices A and S, and the vector b represent the parametrization of the forward model that can be adapted on-line by a supervised gradient procedure as

$$\Delta A = \eta \xi y^\top,\ \Delta S = \eta \xi x^\top,\ \Delta b = \eta \xi \tag{7.7}$$

with $\xi_t = x_t - \psi(x_{t-1})$. In the applications, the learning rate η may be large such that the low complexity of the model is compensated by a fast adaptation process.

The learning rule for the controller, given in (Martius et al. 2013), was derived from maximization of the predictive information. Based on that, we postulate a new unsupervised learning rule (ULR) which will not be derived here but considered as intuitively grounded by the discussion in Sect. 7.4.2 below. The rule is written as (all quantities are at time t)

$$\frac{1}{\varepsilon}\Delta C_{ij} = \delta y_i \delta x_j - \gamma_i y_i x_j \tag{7.8}$$

$$\frac{1}{\varepsilon}\Delta h_i = -\gamma_i y_i \tag{7.9}$$

where δx_t originally was the prediction error but in the new rule we are free to consider δx as any perturbation of the sensor dynamics. In the experiments described below we used the change of the sensor values in one time step, i. e. $\delta x_t = x_t - x_{t-1}$. δy_t is obtained by backpropagating δx_{t+1} through the world model as

$$\delta y_t = J^T \delta x_{t+1} \tag{7.10}$$

where

$$J = \frac{\partial \phi(x,y)}{\partial y}$$

is the Jacobian matrix of the model relating its output to the input y. In our linear model, we simply have $J = AG'(z)$ where $G'(z) = diag[g'_1 \ldots g'_m]$. Moreover, γ_i is a neuron specific learning rate defined as

$$\gamma_i = 2\beta \delta y_i \delta z_i \tag{7.11}$$

where β is the so called sensitivity, an empirical quantity with $\beta > 1$, and

$$\delta z = C\delta x$$

is the perturbation of the post-synaptic potential due to δx_t.

7.4.2 Anti-Hebbian and Differential Hebbian Learning: A Productive Competition

The specific form of the learning rule allows for a very basic interpretation. Let us start with the last term $-\gamma_i y_i x_j$ contributing to ΔC_{ij} which is easily recognized as a Hebbian like term since it is the product of the input x_j into the synapse C_{ij} times the activation y_i of the neuron. In standard situations, the prefactor γ_i is positive, so that overall the term represents an anti-Hebbian mechanism. As such it would weaken all tracks in the SM loop for which there is a strong output of the motor

neuron combined with a strong response from the outside world as reported by the sensor value x_i. In general, this would weaken the activation of any neuron in such a loop, preventing it from saturation.

The first contribution given by $\delta y_i \delta x_j$ is formulated not in the excitations itself but in their time derivatives. This is a differential Hebbian mechanism as it has found much interest recently, see (Kolodziejski et al. 2008, 2009; Kulvicius et al. 2010). Given the relation between δy_t and δx_{t+1}, see Eq. (7.10), C_{ij} is strengthened if there is a strong correlation between[3] δx_j^t and those components of δx_{t+1} which are fed by δy_i^t. Roughly speaking, the first term in the learning rule tries to increase the dynamical correlations across time, driving the system towards activity, while the second term keeps the neurons in their sensitive regions, away from saturation.

This has an important impact on the symmetry breaking scenarios. In the bifurcation scenario discussed in Sect. 7.2.3, symmetry breaking was induced by changing the controller parameter from outside: when crossing the bifurcation point, the system became bistable, the system state jumping into one of the two emerging alternatives. What we have now is a self-referential system, a dynamical system that changes its parameters by itself (Der and Martius 2012). The decisive point in this scenario is the fact that (i) the learning rule does not introduce **explicitly** any violations of symmetries of the physical system it is applied to, but that (ii) the learning is driving the physical dynamics towards activity, eventually causing a **spontaneous** breaking of existing symmetries. This paper will present many examples of behaviors arising from this mechanism of spontaneous symmetry breaking.

7.4.3 Relation to Infomax Principles

As mentioned already in (Martius et al. 2013), there is a close relationship of the learning rules to the so called Infomax principles. Maximizing the mutual information between input and output of a unit is widely known as InfoMax. Applied to a layer of neurons the principle yields an explicit parameter dynamics (Bell and Sejnowski 1995) structurally similar to the one presented in (Martius et al. 2013) and the one given here. Also, similar rules have been obtained in (Triesch 2005) where the entropy of the output of a neuron was maximized under the condition of a fixed average output firing-rate. The resulting dynamics is called intrinsic plasticity as it acts on the membrane instead of on the synaptic level and it was shown to result in the emergence of complex dynamical phenomena (Butko and Triesch 2005; Lazar et al. 2006; Triesch 2007; Lazar et al. 2011). In (Markovic and Gros 2010, 2012) a related dynamics is obtained at the synaptic level of a feedback circuit realized by an autaptic (self) connection. In a recurrent network of such neurons it was shown that any finite update rate (ε in our case) destroys all attractors, leading to intermittently bursting behavior and self-organized chaos.

Our work differs in two aspects. On the one hand, while the information theoretical principle—the starting point for our more intuitively based learning rule—was formulated at the level of behaviors of the whole system, explicit learning rules were

[3] We write the time as a superscript if components of a vector are considered.

derived at the neuronal level by rooting the behavioral level (outside world) back to the level of the neurons (internal world). On the other hand, as a direct consequence of that approach, there is no need to specify the average output activity of the neurons as in (Triesch 2005). Instead, the latter is self-regulating by the closed loop setting. Independent of the specific setting, the general message is that these self-regulating neurons realize a specific working regime where they are both active and sensitive to influences of their environment. Instead of studying those neurons in internal (inside the "brain") recurrences, we embed them into a feedback loop with complex physical systems where these self-active and highly responsive neurons produce surprising phenomena at the level of behaviors—in the outside world.

7.5 Homeokinesis: Body Inspired Behavior

The unsupervised learning rule (ULR) given above shares some common features with the corresponding rule forming the basis of homeokinesis. In the Chapter by Martius, Der, and Hermann, the latter approach will be considered in some detail, demonstrating the potential of that approach for guided SO. This section will give a very short account of the principle of homeokinesis in order to discuss the parallels and differences of the two related approaches. Reader not interested in this comparison may skip this section and go directly to the applications.

Homeokinesis (HK), a general principle originally proposed in (Der 2001; Der and Liebscher 2002), stands for the dynamical symbiosis between brain, body, and environment. It was shown by many works to drive robots to a self-determined, individual development in a playful and obviously embodiment-related way, see (Der and Martius 2012) for a detailed consideration. HK is not only the description of a goal, namely the dynamical symbiosis between brain, body, and environment but comes with a concrete realization in the form of a universal, unsupervised learning rule.

In order to discuss similarities and differences with the ULR of this work, we directly consider the HK rule obtained by the minimization of the so-called time loop error (TLE)[4]

$$E_{TLE} = v^T v = \zeta^T \frac{1}{CC^T} \zeta \; . \tag{7.12}$$

where ξ is the prediction error[5], $\zeta = (AG')^{-1} \xi$, $\mu = \left(CC^T\right)^{-1} \zeta$, and $v = C\mu$. Gradient descending the TLE leads to the following unsupervised learning rules for the parameters of the controller

[4] The matrix inverses have to be understood as pseudo inverses if the normal inverse does not exist.

[5] The minimization of the prediction error is also at the basis of more recent approaches on minimizing the free energy of the sensor process (Friston 2010; Friston et al. 2012; Friston and Stephan 2007; Friston 2012). In Friston's approach, the tendency of stasis is overcome by assuming additional priors that drive the system to activities. This is not necessary both in homeokinesis and in the present approach.

$$\Delta C_{ij} = \varepsilon_c \mu_i v_j - \gamma_i y_i x_j \,,$$
$$\Delta h_i = -\gamma_i y_i \tag{7.13}$$

displaying a noteworthy similarity with Eqs. (7.8, 7.9). Indices are running over all sensors and motors as before, i. e. $i = 1,2,\ldots,n$ and $j = 1,2,\ldots,m$ and we introduced the channel specific learning rate

$$\gamma_i = 2\varepsilon_c \mu_i \zeta_i \,.$$

These relatively simple update rules define the parameter dynamics of the controller, the learning of the self model being given by Eq. (7.7). The rules need some numerical precautions due to the matrix inverses which are discussed in detail in (Der and Martius 2012). Both in HK and in the present work, learning is not to be understood as the convergence towards a specific goal. Instead, the learning rates usually are chosen such that the parameter and system dynamics run on comparable time scales. In the neural network interpretation we have a fast synaptic dynamics which is constitutive for the behavior of the system.

7.5.1 Principles of Action

The peculiarities of the general learning rules has been demonstrated in various robot examples using both real and simulated robots, see (Der and Martius 2012). At a general level, essential features of a learning rule are revealed by considering the landscape of the objective over the parameter space. The time loop error, Eq. (7.12) is characterized by singularities acting as repeller (infinitely repulsive regions) for the gradient flow. This is a direct result of learning the system backward in time. These repulsive regions are easily identified. A first class of singularities is given by the zeros of g' (featuring in ζ), i. e. in the saturation regions of the neurons. The effect of this singularity, which leads to the anti-Hebbian term $-2\varepsilon_i y_i x_j$ in Eq. (7.13), is to keep the neurons away from the saturation regime. This is very reasonable since in that region neurons are not sensitive to their inputs.

Typical for the landscape based on the TLE is a further repulsive region resulting from the inverted matrix CC^T. The role of this singularity is most immediately seen if the robot is initialized in the situation $C = 0$, $h = 0$, (the least biased one in the sense explained in Sect. 7.6.1) so that the controller does not react to its sensor values at all. The slightest perturbation will quickly drive the parameters of the controller away from this unstable fixed point of the combined dynamics. In this way, feedback in the sensorimotor loop is generated and the robot is quickly driven away from this "do nothing" region in behavior space. This is a definite advantage if it is important to get away from that singularity as fast as possible. In our new rule, the "do nothing" region is not a repeller but an unstable fixed point. This may be a disadvantage in the bootstrapping process. However, with the new rule, one may start directly with the least biased initialization. This is beneficial if one wants to make maximum use of the symmetry breaking phenomenon. This paper presents several examples which were not possible in the HK approach.

Another singularity is produced by the inverse of the self-model matrix A featuring in ζ. That one can also be present in the new approach if δy is calculated not by backpropagation but by the so called backprojection method. We consider this as a more or less technical issue so that we will not discuss it here, see (Der and Martius 2012) for details.

Behaviors generated by either the HK or our new rule are inherently contingent (there is no influence from outside) but by far not arbitrary since the whole bootstrapping process is driven by the specific reactions of the embodied robot to the controller signals. Thus, it is the body itself which plays the most active part in the emerging control process so that this phenomenon has also been called body-inspired behavior.

7.6 Vehicles: Behavior As Broken Symmetry

Let us now apply the new learning rule to some specific examples chosen such that the characteristic properties of the self-organization process are illustrated. This section will focus on the TWOWHEELED robot.

7.6.1 Least Biased Initialization

In applications, a first point is about the choice of the initial parameters of the networks and the initial configuration of the robot. With our specific choice of the controller network, the initialization with $C = 0$ seems most natural because this corresponds to a controller that is completely numb, i. e. deprived of any functionality. Putting additionally $h = 0$, we find that all motor neurons send the command $y_i = 0$ to the controller, independently of any inputs.

Choosing the initialization in the described way has different effects on the initial pose the robot is taking. For example, in the wheeled robots case this means that all wheels are held at rest. In robots consisting of several segments tied together by joints, $y = 0$ means that all joints are driven towards their center position. We will follow up this point in the detailed examples of the HEXAPOD and the HUMANOID further below.

The combined system, comprising the physical and the synaptic dynamics, is fully deterministic, in the virtual case at least. If starting in the least biased initialization the combined system may be in an unstable fixed point so that we have to add, for a short time interval, a little noise to the sensors. After switching the noise off, the actual initial condition is fixed. From this time on, the further time evolution of the entire system is deterministic, being fully determined by the values of the sensor vector, the current pose of the robot and the values of all parameters of the controller (now already different from zero). This set of values may be called the behavior code of the robot.

7.6.2 Symmetry Breaking—A Rule of Thumb

Before going to present the experiments, let us formulate a simple rule of thumb on the development of the robot when starting from the least biased condition: in typical experiments we observe that the behavior of the robot can be described as being active (caused by the differential Hebbian term in the learning rule) while conserving as much of the original symmetries of the system as possible. This was also formulated as the economy or parsimony of symmetry breaking. Note that symmetries involve not only the geometry of the robots body but also all the symmetries of the physical dynamics. In the two-wheel robot case the body geometry is described by the invariance against both left-right and forward-backward transformations. The physical symmetries are based on the robot being an object in space and time, the physics being invariant against both translations and rotations of the frame of reference, taking however account of the physical boundaries (objects, the ground the robot is on). Moreover, there is the invariance against the inversion of time combined with the inversion of all velocities.

Let us also emphasize that symmetry breaking observed in this scenario is emerging as a phenomenon "from inside" the deterministic system itself so that we may speak of a spontaneous symmetry breaking (SSB). We would also like to stress that fundamental modes are singled out by a maximum degree of symmetry. This may give the concept of fundamental modes a more solid founding but this still needs some work to be done. As an additional feature, the breaking of the symmetries can largely be influenced by external impacts (physical forces in the sense of a desired mode) and/or by choosing specific sensor combinations that help to organize the symmetry breaking scenario. We will give an example with the bungee jumper below.

7.6.3 The Autistic Vehicle: Fundamental Modes

We have introduced above the vehicle with the wheel counters as the only sensors. We have called it autistic because its only contact with the external world is by proprioception. Moreover, the learning starts in the least biased way, so that the symmetry breaking should follow the principle of parsimony mentioned above. In particular, the physical system (realized in the simulator by N differential equations, where $N = 8$ with the vehicle), is invariant against spatial transformations, i. e. translations or rotations of the spatial frame of reference. With the constraints given by the (elastic) surface, the remaining symmetry operations are rotations around the z axis and translations in the $x - y$ plane. Also, as discussed, the learning rule gives no clue of how symmetries are to be broken.

We consider now the TWOWHEELED with a two-neuron controller receiving sensor vector $x \in \mathbb{R}^2$ representing the wheel velocities and sending its output $y \in \mathbb{R}^2$ to the motors. Starting with the least biased initialization, and believing in our rule of

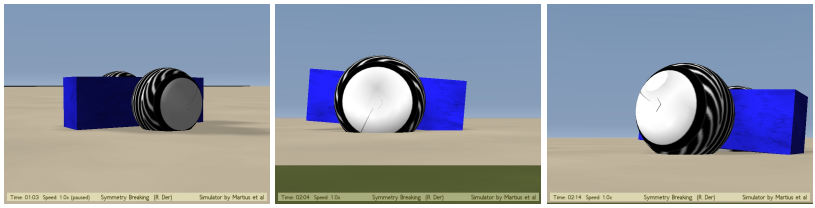

Fig. 7.2 The TWOWHEELED as a 3D physical object. The ground is elastic so that the wheels are sinking in, depending on their load. Left, wheelsize = 1: with the given elasticity, the robot is lying more or less flat on the ground when driving straight. There is a strong slip effect when accelerating. Moreover, when moving in a curve, there is an inclination due to the physical forces making the effective radius of the wheels different, see the video S2. This is even more pronounced with larger wheel sizes (middle and left, size = 1.2). These 3D effects make both odometry and the execution of motion plans very difficult as they involve the full physics of the robot.

thumb, we expect the robot to start moving after some time[6] while trying to conserve as much of the original symmetries as possible.

However, when using a learning rate ε not too small, the robot is seen to engage in a sequence of left and right turns combined with motions back and forth along curved lines, without any regularity to be seen, see Fig. 7.3, top left panel. Still, we must note that these trajectories are fully deterministic. Restarting the simulation with exactly the same parameters C, h and in the same physical state reproduces exactly the same trajectory however irregular it is. Nevertheless, our rule of thumb obviously is not valid in this regime as there is no visible footprint of the underlying symmetries—the invariance against rotations and translations of the physical space.

The situation changes drastically when using smaller learning rates so that the interplay between learning and physical state dynamics is given time to unfold. Fig. 7.3 is demonstrating a typical behavior of the robot. As the top right panel shows, after starting, the robot is running through a kind of metastable patterns, converging after some time toward a large scale CP. The bottom left panel shows the fine structure of that limit cycle behavior.

In order to discuss this phenomenon, let us take at first the perspective of an external observer in the sense of Braitenberg. From this external point of view, the robot is drawing a complex geometrical pattern with a surprisingly high accuracy. In order to do that, the robot would have to have a detailed plan of the pattern to be drawn. In order to execute that plan, the robot would need a very accurate system of self-localization for finding its exact local position and heading direction. This could be done by either using a compass and a GPS system (which is not available) or by odometry, based on the measured wheel velocities. However, this is very inaccurate, given that the robot in the simulation is moving on an elastic

[6] When using low learning rates, this time can be very long so that we often start the robot with an initialization close to the bifurcation point, choosing $C = c\mathbb{I}$ with c close to 1. Contrary to the HUMANOID and HEXAPOD treated below, in the TWOWHEELED case, no substantial differences in the behaviors were observed.

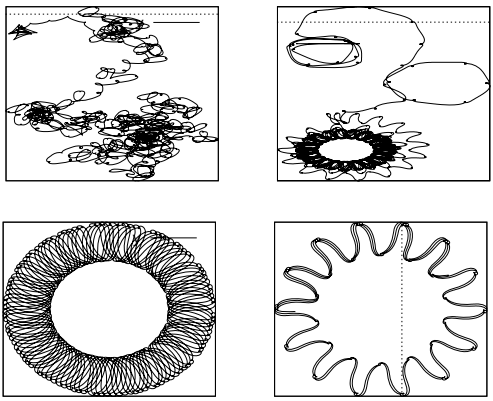

Fig. 7.3 Deterministic trajectories of the robot in the 2D plane emerging with different learning rates ε. If learning is fast ($\varepsilon > 0.01$), irregular trajectories are the rule. With lower rates (here $\varepsilon = 0.001$), after a transient phase of irregular motion through a kind of metastable attractors, the dynamics is converging toward a limit cycle behavior (top right and bottom left panel). A first explanation for the emergence of a circular pattern (CP) is given by using an SO(2) matrix for controlling the robot (bottom right). This leads to a CP but does not explain the fine structure of the CPs observed under learning.

ground so that, when accelerating or moving in a curve, there is an inclination due to the inertia and/or centrifugal forces making the effective radius of the wheels different, see Fig. 7.2. On top of that, the odometry would be very demanding in the computational resources given the required accuracy.

7.6.4 Synergy of Learning and Physical State Dynamics

On the other hand, when looking inside the system, none of the above ingredients for drawing a prescribed pattern with the observed accuracy can be found. Instead, all we see are the two poor neurons with their universal, low-complexity synaptic dynamics. Metaphorically, from this point of view, nothing like a concept of space and/or a plan for drawing a particular pattern can be uncovered. Even more so, it is even impossible of finding the plan inside the robot. In fact, given the wiring, the emerging patterns are specific with the concrete embodiment as they emerge by the interplay with the physics of the system only.

Still, an understanding of the phenomenon is in reach when looking at the structure of the controller in one of the limit cycle patterns. When taking a snapshot of the controller matrix C, we discover a nearly perfect $SO(2)$ structure of the C-matrix, see Fig. 7.4, with tiny modulations in synchrony with the motions of the robot. Any such matrix can be written as a rotation of a given vector v by an angle α and a stretching factor ρ. When using such a matrix in Eq. (7.5), the robot will either draw a (nearly) perfect cycle or, as shown in Fig. 7.3 (bottom right), will generate a circular pattern similar to the co-learning case but with a very different fine

structure. Obviously, the fine structure of the patterns is produced by the coupling between physical and learning dynamics so that the rotation matrix does give only a first orientation on the emerging structures. Instead, as we will see below, both the very nature and the fine structure depend additionally on the embodiment in an intricate way.

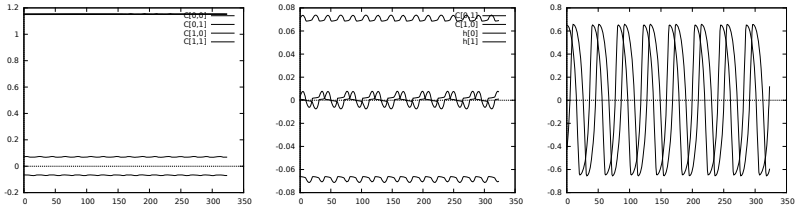

Fig. 7.4 The circular pattern formation is hidden in the dynamics of the controller parameters as driven by the general learning rule. The C-matrix (left and middle panel) is seen to be of a nearly perfect $SO(2)$ structure. This reveals another aspect of the general learning rules—the emerging dimension reduction in the parameter space. Instead of 4 parameters, the $SO(2)$ matrix is described by only one parameter (the rotation angle). The middle panel displays the interplay with the h dynamics which is seen to be periodic with a slight bias. Right panel: the two sensor values (wheel velocities).

So, from looking inside, there is a coarse explanation of how the robot does it, once it has "understood" so to say, what a rotation matrix is and how it can be used. However, this does not explain why the learning discovers these special matrices in the four dimensional matrix space and moreover how to use them in the fine interplay between the parameter and the (physical) state dynamics.

On a general level, an understanding is given by or rule of thumb: a pattern in space can only emerge from breaking the spatial symmetries inherent in the physics of the robot as mentioned above. When trying to make this symmetry breaking as parsimonious as possible, a circle is nearly perfect: while it has broken the translational symmetry (the center is a fixed point in space), rotation symmetry (around that center) is fully conserved. Yet, because of its fine structure, the real patterns emerging in our learning scenario are not circles, but they are still invariant against rotations about a definite angle, see in particular the patterns of Fig. 7.5. This may be seen as a noteworthy parallel to the hexagonal patterns known from many phenomena in nature. So, the observed patterns apparently are the ones with a high degree of preserving the spatial symmetries of the physical system.

There is a more naive argument for the predominance of such regular patterns. We have seen above that, once the initial condition is fixed (see Sect. 7.6.1), the trajectory of the robot is exactly determined for all future by the values of that initialization—its behavior code. Our point is that a completely irregular trajectory needs infinitely much information stored in its code. To the contrary, codes of regular trajectories need much less information so that there are whole sets of values generating one and the same regular trajectory (in a finite but large time interval) so that the latter are realized with much higher probability.

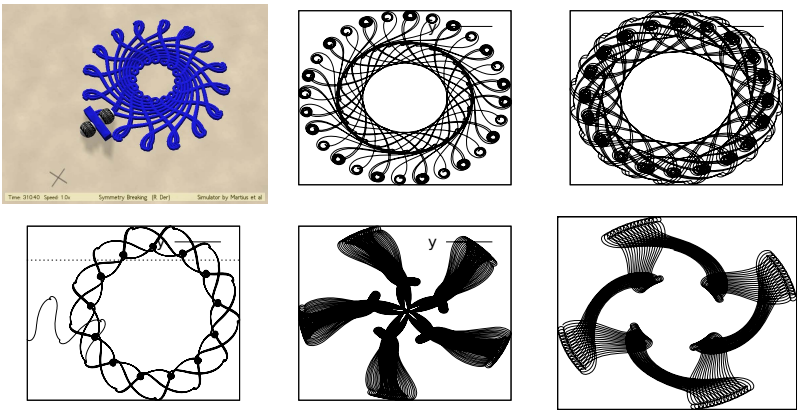

Fig. 7.5 Pattern spin-off effect in the learning on-off scenario. When in a limit cycle (circular pattern) and temporarily switching learning off, a new stable structure is emerging. After switching learning on again, the full dynamics is converging back to the limit cycle. In repetition, this produces a large variety of different structures. Top left: Pattern with robot for size comparison. However, patterns come on very different length scales. See video S3 for an example.

7.6.5 The Pattern Factory

Further insight into the role of the learning dynamics in the pattern formation can be gained in a kind of learning on-off scenario: we start with the full learning dynamics, let a circular pattern, see Fig. 7.3, develop and put $\varepsilon = 0$ for a while. Most astonishingly, the robot starts drawing a new pattern of very high geometric regularity resembling only remotely the circular structure it was spinning off. Upon switching on learning again, the system rapidly returns to the original CP (with a different center in space). This procedure can be repeated with ever new patterns emerging.

In discussing this pattern spin-off effect, we have to remember that once $\varepsilon = 0$, both C and h are fixed so that the spin-off pattern is fully determined by the physical starting state. In a number of experiments we even observed that each parameter set created in the learning on-off scenario is seen to even support several behaviors, each with its own basin of attraction. What we want to emphasize is that all these patterns are off-springs of the CP with behavior codes defined by the snapshots of the parameter dynamics. So, actually, we may say that, once in (or close to) the limit cycle attractor, the learning (parameter dynamics) produces a whole sequence of behavior codes for potential behaviors.

7.6.6 Patterns As Expressions of Embodiment

Let's now have a closer look at the role of the embodiment. As already explained with Fig. 7.2, the TWOWHEELED must be considered as a full 3D system moving on

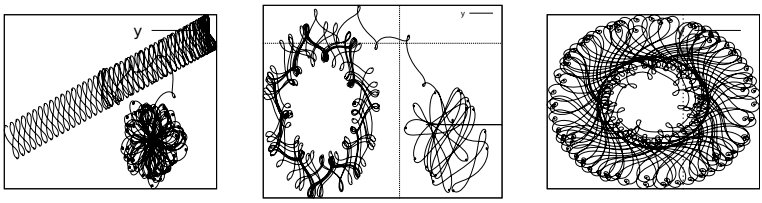

Fig. 7.6 The role of embodiment. Left panel: wheel size 1.1, $\varepsilon = .001$ and cInit=1.2. After a very irregular initial phase, the robot enters an aligned wiggly pattern, running at first to the right and then back toward the left lower corner. Middle: Wheel size 1.15. Right: asymmetry of 1.5:1.75 in the wheel-trunk system with wheel size 1.2. The pattern is less regular.

an elastic ground. Let's consider a few experiments demonstrating how sensitively the emerging behaviors depend on the concrete realization of the robot's body.

As explained with Fig. 7.2, the physics of the robot already changes substantially with changing the wheel size. As demonstrated with the left panel of Fig. 7.6, changing the wheel size from 1 to 1.1 leads to a new pattern—instead of a circular structure we get convergence toward an aligned pattern structure, i. e. a wiggly structure organized along a line. On the general level we may argue that now it is not the rotation symmetry that is partially conserved, but the translational one. In fact, a line is invariant against translations but not rotations in the plane.

As the middle panel of Fig. 7.6 shows, this effect is very sensitive to the choice of the wheel size—increasing the size to 1.15 produces a circular pattern again but with a very different fine structure. The emerging patterns are also very sensitive to the forward-backward asymmetry of the wheel-trunk system, see right panel of Fig. 7.6 for an example. Depending on that degree of asymmetry many different patterns can be observed, mostly circular ones but with widely varying fine structure.

7.6.7 Modes

We see that the hidden symmetries of the physical system can in a very intricate way propagate into the combined state-parameter dynamics of the developing system to produce large-scale effects in space and time[7]. When looking at the limit cycle patterns (the circular structures), we see a very complicated periodic or quasi-periodic structure interweaving the physics and the learning dynamics in an irreducible way. This synergy between state and parameter dynamics is breaking down

[7] This effect seems to have been observed already in the homeokinesis (HK) approach (Der and Martius 2012) with the sphere that adapts its rolling motion to the geometry of the basin. Here, we have a more subtle effect, given that there is no formative external geometry, the only link between the internal world (learning dynamics) and outer world (physics) is by the wheel velocities which give information about the behavior in space only in an indirect way and, even in the simulations, with much restricted accuracy. It is only by this window into the outer world by which physical effects, mainly the inertia due to the mass of the body and the wheels, take influence on the behavior generation.

immediately as soon as we freeze the learning dynamics, with a new pattern emerging as explained above.

We may consider each emerging limit-cycle dynamics as a mode of the system. These are of different nature, either with the parameters fixed or driven by the learning dynamics. In the no-learning case, there is an immediate parallel to Braitenberg's idea of interchanging perspectives: looking from inside, we see a vehicle with a fixed wiring of sensors to motors, now also with cross connections and some bias for the neurons (active wires). This extremely simple electro-mechanical system is able of drawing very complex geometrical patterns like the ones shown in Fig. 7.5. Looking from outside, even if the self-localization problem was solved, the system seems to follow a very complex plan for drawing such highly regular geometric patterns. But also in the learning case, the Braitenberg perspectives are still appropriate. If considering the parameter dynamics as the dynamics of some internal state variables, we still have a very simple controller exciting those very complex modes.

Yet another perspective is suggested by the autism metaphor: we could say that, by the learning dynamics, the robot is somehow taking note of its being in space (creating patterns) but this without any internal reflection of that fact. Metaphorically, we may say that the robot does not know what it does but it does this with high accuracy and dedication, discovering by proprioception alone a specific way of being in space and time.

The observed limit cycle attractors, both with and without learning, are examples of the many modes this specific brain-body system is able of developing. We are free to call each of them a fundamental/normal mode but we also may single out the co-learning mode as the most fundamental one as it is the host of all the spin-off modes. This is the attitude taken in this paper.

7.7 The Looping HUMANOID

Let us now continue our considerations with a more complex system, the HUMANOID introduced above. The results in Sect. 7.3 have demonstrated that under certain conditions this system is particularly amenable to fundamental modes. The observed modes may be attributed to a specific form of broken symmetries under minimalistic control. Using the HUMANOID we are now going to look a little deeper into the scenario of spontaneous symmetry breaking in systems of higher complexity than that of the TWOWHEELED.

7.7.1 High Symmetry Motion Patterns

As in the TWOWHEELED case we start with the least biased initialization, wondering what the emerging symmetry breaking scenario will look like. We use the bungee scenario with the same physical setting as in Sect. 7.3. After letting the robot fall, we observe again the vertical oscillations around the equilibrium point but after a while (several minutes real time with $\varepsilon = .001$) the robot develops a specific motion pattern looking like a swimming motion, see Fig. 7.7 and video S4.

On a general level, the observed pattern illustrates very clearly the rule of thumb: the robot is being active but the original symmetries of the physical system are conserved as much as possible. This is obvious for the geometrical left-right symmetry but we may also argue that the original symmetries of the physical system are still dominating the behavior. In the least biased initialization, there is an invariance against inversion of joint angles[8]. Breaking the symmetry means to define for each joint angle where to go by what velocity. Without any external clue, the most parsimonious solution would be to change all the joint angles with the same velocity, generating a phase and frequency locked dynamics of the mechanical system. Actually, this is what we observe at least in the early phases of the behavioral self-organization.

Looking more into the details, we see a qualitative difference to the running motion pattern (RMP) of Sect. 7.3 with its in-phase and anti-phase signature. As all other conditions are equivalent, this must be caused by the different initializations. In the RMP, the strong negative feed-back strength of the system was discussed responsible for the specifics of the RMP, see Sect. 7.3.2. In this setting, we start with $C = 0$ and $h = 0$, the least biased initialization, everything being determined by the learning, starting with the initial perturbations introduced by the physics. In the beginning, the vertical oscillations introduce tiny changes in the joint angles due to inertia effects not fully counteracted by the motors. These are acting on all joints simultaneously as they are caused by the motions of the center of gravity. Once started, these perturbations are self-amplifying which may explain why all angles are changing simultaneously. Later on, there may also be some physical cross-talk as described with the RMP in Sect. 7.3. This is now possible not only in the joints of a single limb but may propagate also between the upper and the lower part of the body and between the arms and legs.

The cross-talk phenomenon shows that the symmetry breaking mechanism is very sensitive to the concrete morphology of the robot. By way of example, we use an experiment where the robot was given an additional back joint, giving the robot more freedom in movements of its lower part. Moreover, we fixed the trunk so that there is no starting perturbation by the vertical oscillations as in the bungee setting. Still, there is a physical starting signal given by the difference of the joint angles in the initial configurations (all equal to zero) and their equilibrium position (different due to gravity). Again, after a transient initial phase, we observe the emergence of a collective mode involving all degrees of freedom of the lower part, see Fig. 7.8. We guess that this collective motion is more a result of the physical cross-talk than the fingerprint of the initial symmetries, but we still have to study this in more detail.

7.7.2 Exterioception May Guide Self-Organization

Up to now we were using only proprioceptive sensors so that the orientation of the bungee jumper in space came into play only in a very indirect way, for instance by

[8] In order to make this explicit, think of a linearization of the system around its equilibrium point defined by the initialization. Details will be given elsewhere.

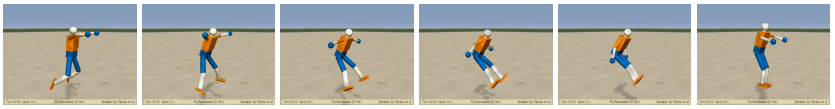

Fig. 7.7 One period of an emerging motion pattern in the bungee scenario. The pattern is an example of minimal symmetry breaking.

Fig. 7.8 One period of an emerging periodic motion pattern with the trunk fixed and an additional back joint giving the robot more freedom of motion. Coherence is now caused not only by the original symmetries but also by the physical cross-talk becoming "louder" with increasing activity of the robot.

the gravity forces providing additional loads on the joint motors. Let us now include a sensor that is orientation sensitive. In a first experiment we use a sensor that measures the rotation velocity of the trunk around its yaw axis (the one running through both the shoulders). This velocity is simply taken as an additional component of the sensor vector x. We let the system run with the learning algorithm of Eq. (7.8) just as in the previous case.

Astonishingly, the emerging behavior is completely different from the previous one. We observe now already after a short time that the robot tries rotational motions around that axis (known to him only by one sensor value without any knowledge what this sensor is measuring). In the video S 5, the amplitude of these rotational motions is seen to increase steadily until the robot is looping, see Fig. 7.9. After that it is in a physically less stable state (by loss of symmetry in its motion pattern) so that it needs some time until the play may begin again.

Again, we can recognize the rule of thumb here. In fact, the motion still—while being of quite some variability—reflects the original symmetries of the physical system in its least biased initialization. In particular, the looping is performed only if the robot is keeping in a motion regime where the left-right symmetry with respect to the yaw axis is nearly perfect, see the video S 5.

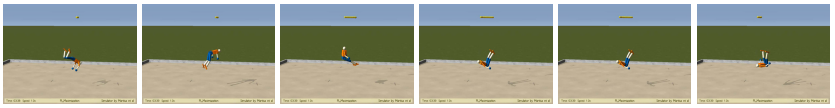

Fig. 7.9 With a sensor measuring the rotation velocity of the trunk around its yaw axis, the robot synchronizes its internal motions with the trunk rotation. After a while, the robot learns rotating its entire body around the trunk axis, eventually executing a loop.

7.7.3 Starting in a Mode

What happens if we start with the above minimalistic controller, i.e. use the initialization $C = c\mathbf{I}$ with $c < -1$? As video S 6 shows, the learning enhances the running motion pattern (RMP) in the course of time more and more by steadily increasing the step length until the whole motion pattern gets unstable after a long time. However, as video S 6 also shows, the stability of the RMP is holding up to a very large step length. This video also shows that the RMP is entirely self-generated. When the bungee force is reduced so that the robot reaches ground, the RMP decays but is resumed rapidly if the conditions of its existence (hanging in air) are reestablished.

7.8 The HEXAPOD

Let us now follow the trace of symmetry breaking a little further down the road with a robot of comparable complexity as the HUMANOID but with a different symmetry and interaction with the environment: the HEXAPOD. We choose this robot since it has not been treated in earlier work on HK (Der and Martius 2012) and because it will be seen to reveal symmetry breaking phenomena in a particular clear way. The robot consists of a trunk with two whiskers (passive joints), and six legs, each one consisting of a shoulder with two joints and a knee with a single joint. Each of the 18 joints is activated by a motor and contains a sensor that is measuring the true joint angle. The effective torques acting on the joint axes are determined by a PID controller so that there is an elastic reaction of the robot to the nominal joint positions $y \in \mathbb{R}^{18}$, similar as in a system controlled by muscles.

In a typical experiment, the HEXAPOD is falling down from a starting position a little above the ground. With the least biased initialization, $y = 0$ in the beginning so that all joints are in their center positions. When hitting the ground, the robot gets into a damped vertical oscillation due to the elasticity of the joint-motor system. This is sufficient for providing an initial perturbation that is further amplified by the bootstrapping mechanism as sketched in Sect. 7.6.

7.8.1 Modes

What can we expect to happen? Depending on the concrete situation (for instance the meta learning parameter, in particular ε) different behaviors may emerge. In most cases the robot starts with a swaying and rolling motion pattern, see Fig. 7.10 and video S 7. We may claim again, that this is in agreement with our rule of thumb since in this motion the joint angles are changing with a pretty high degree of coherence as allowed by the (soft) physical constraints enforced by the ground contacts.

More interesting behaviors emerge after some time. Most exciting is the emergence of a jumping behavior as a stable phenomenon, see Fig. 7.11 and the corresponding video S 8 and S 9. In these jumping patterns the robot is seen to be very active but still with a large coherence in the joint angles as suggested by our rule of thumb. There is still a surprise when looking at the parameters of the controller.

Fig. 7.10 Initially, the robot develops a swaying motion pattern, as if it is very actively trying to move the legs in a coherent way while keeping ground contact

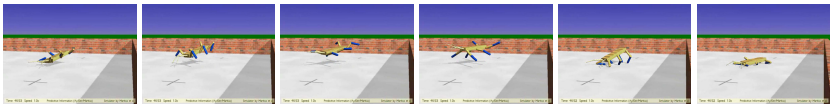

Fig. 7.11 Emerging jumping motion pattern

In the TWOWHEELED case we observed that the C-matrix displays a definite structure, it developed into a $SO(2)$ matrix which is prominent among all $\mathbb{R}^2 \times \mathbb{R}^2$ matrices in forming a group. This is like the symmetries of the physical space leaving a fingerprint since the physical space in $2D$ is invariant against all transformations by the group of $SO(2)$ matrices.

Of course, we can not expect such a clear result in the case of our HEXAPOD because of the much higher dimensionality of the physical space and the fact that the robot is also interacting with the ground. Yet, a look at the C-matrix is the more surprising. As Fig. 7.12 shows, the emerging sensor-to-motor coupling matrix is highly structured, reflecting the original symmetries to an amazingly high degree. Both the shoulders and the knees are seen to follow essentially the same strategy for

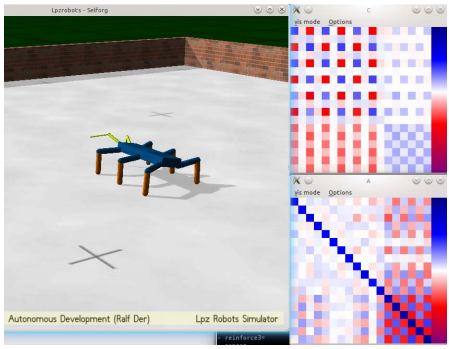

Fig. 7.12 The parameters of the controller (C-matrix, upper right panel) and the model (A-matrix). In the C-matrix, the upper left 12×12 block displays the matrix elements of the shoulder couplings, each row i representing the coupling strength of the 12 angle sensors into the i^{th} motor neuron. Reading each row from the diagonal element to the right and wrapping the reading back to the diagonal, one sees that each shoulder follows essentially the same strategy. A related understanding of the control strategy holds for the knee joints, see the lower right sub-matrix.

moving the body. This is in agreement with our rule of thumb since this collective strategy allows the body to be moving, even to jump, but with a maximum degree of coherence between the individual constituents of the body.

Moreover, we also see from that figure the whole-body nature of the behavior. The periodic jumping pattern are not created by using a central pattern generator producing a master signal that is sent with the necessary phase corrections to the individual motors. Instead, the motion of each body part is generated by combining both excitatory and inhibitory signals from the sensors of all joints in a systematic manner. This is a new control strategy for generating jumping (and, hopefully, also walking like) patterns emerging from simply applying the general learning rule to the physical system.

7.8.2 Perspectives for Guidance and Reinforcement Learning

Important is also the observation that, by including exterioceptive sensors, the development of the modes can be influenced and driven into desired directions. For instance, as our experiments show (to be presented in detail in a later paper), the jumping height can be increased by simply including a sensor measuring the vertical velocity of the robot trunk. The effect can be increased if, like in reinforcement learning, the parameter updates are weighted according to the agreement with the reward. For instance, by using the forward velocity of the robot as a reward, we can guide the SO process towards jumping into the forward direction, see Fig. 7.13 and video S 10 as an example.

There are many more interesting phenomena observed with that kind of robots but we are not going into more details here. We think that the few examples given may serve as an outline for the bunch of phenomena emerging from SSB under the general learning rule, given a convenient initialization and some patience for doing the experiments.

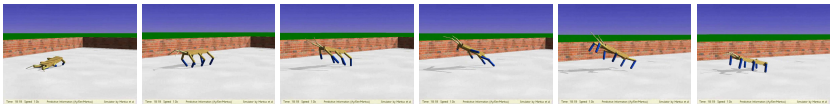

Fig. 7.13 Jumping motion pattern driven by rewarding the forward direction

7.9 Concluding Remarks

This paper tried to answer essentially two questions. The first question is about how to organize self-organization, asking how can we find intrinsic mechanisms that make a system able to self-organize. The answer was given by the unsupervised learning rule (ULR), see Eqs. (7.8, 7.9), that drives systems into self-organization. The ULR stages a weighted competition between a differential Hebbian and an anti-Hebbian learning mechanism. While the former drives the system to activity, the latter acts as a confinement, keeping the system under control. The rule

is strictly local—synaptic changes exclusively depend on excitations at the ports of the synapse—yet creates global, whole-body motion patterns of the robotic system. This is demonstrated in applications with (i) a wheeled robot that spontaneously is moving in geometric patterns, and (ii) a hexapod robot with 18 degrees of freedom self-organizing from scratch into several locomotion patterns like jumping and walking.

The second bunch of questions is suggested by exactly that bootstrapping scenario: with nothing specified from outside, what can we expect the learning system to do. What will the emerging behaviors look like and what will the relation to the embodiment of the robot be? How and to what extend are the emerging behaviors determined by the embodiment; and can we find systematic criteria for those behaviors?

Several answers could be given by looking into the role of the underlying symmetries of the system in space and time. The invariance of the physical space against symmetry operations (like translations and rotations) induces corresponding invariances modified by the constraints (like contact with the ground). The point then is that, while driving the system towards instability, the ULR is preserving these symmetries. As a result, the evolution of the system in the learning process is realized by a sequence of spontaneous symmetry breaking following—similar to what we know from nature—a kind of parsimony principle. This lead to our rule of thumb: the emerging behaviors in physical systems (robots) driven by our ULR are qualified by a high activity while preserving as much of the underlying symmetries as possible.

This rule brings the embodiment to the foreground. The symmetries are embodiment specific and, moreover, breaking the symmetries is a process that is related to the very physics of the system. This was demonstrated by a number of examples. The first and probably the most surprising one was given by the TwoWheeled robot. Controlled by two neurons with a fast synaptic dynamics given by the ULR, the system in many cases was converging towards a limit cycle behavior with the trajectories of the robot forming nearly perfect geometric patterns. The emerging geometric patterns where seen to depend on the embodiment (like the wheel size) in a very intricate and sensitive way. Interestingly, the limit cycle features as a pattern factory: in a learning on-off scenario it was seen to produce a great variety of further patterns, spinning off the limit cycle if the learning is turned down (called the pattern spin-off effect). It would be very interesting to ground this phenomenon in dynamical system theory.

Similar effects of symmetry breaking were obtained by the examples of the Humanoid and the Hexapod. In the latter case, we also observed the excitation of body related, high activity modes with a high degree of coherence between the body parts. In particular, with the Hexapod, we observed a jumping mode demonstrating our rule of thumb: emerging behaviors are qualified by high activity while preserving the underlying symmetries of the system as far as possible (the principle of parsimony in spontaneous symmetry breaking). In future work we will be looking for a parallel of the pattern spin-off effect, hoping to thereby uncover a kind of pattern factory for these more complex systems, too.

These results are a step forward as compared to the state of the art. Previous work in self-organizing robot behavior was either restricted to small, easy to analyze systems or produced—like with the principle of homeokinesis—behaviors which looked interesting and were often completely surprising (Der and Martius 2012), as it should be. However, by the same token, it was often not clear what the robot actually is doing. With the concept of behaviors as broken symmetries, this is now (a little) different. The principles and examples given in this paper—in particular the emergence of fundamental modes, the TwoWheeled as a pattern factory, the looping behavior of the Humanoid, and last but not least the jumping modes of the Hexapod—may help us to better understand and exploit the synergy between embodiment and SO of autonomous robots.

Acknowledgements. The author gratefully acknowledges the great hospitality in the group of Nihat Ay at the Max Planck Institute for Mathematics in the Sciences and many helpful and clarifying discussions with Nihat Ay, Georg Martius, and Keyan Zahedi. I am also grateful to Mikhail Prokopenko (CSIRO Sydney) for very stimulating discussions on self-organization.

References

Anthony, T., Polani, D., Nehaniv, C.L.: Impoverished empowerment: 'Meaningful' action sequence generation through bandwidth limitation. In: Kampis, G., Karsai, I., Szathmáry, E. (eds.) ECAL 2009, Part II. LNCS, vol. 5778, pp. 294–301. Springer, Heidelberg (2011)

Ay, N., Bernigau, H., Der, R., Prokopenko, M.: Information driven self-organization: The dynamical systems approach to autonomous robot behavior. Theory Biosci. (2012)

Ay, N., Bertschinger, N., Der, R., Güttler, F., Olbrich, E.: Predictive information and explorative behavior of autonomous robots. The European Physical Journal B - Condensed Matter and Complex Systems 63(3), 329–339 (2008)

Behnisch, M., Haschke, R., Ritter, H., Gienger, M., Humanoids: Deformable trees - exploiting local obstacle avoidance. In: Humanoids, pp. 658–663 (2011)

Bell, A.J., Sejnowski, T.J.: An information-maximisation approach to blind separation and blind deconvolution. Neural Computation 7, 1129–1159 (1995)

Braitenberg, V.: Vehicles: Experiments in Synthetic Psychology. MIT Press (1984)

Butko, N.J., Triesch, J.: Exploring the role of intrinsic plasticity for the learning of sensory representations. In: ESANN 2006 Proceedings - 14th European Symposium on Artificial Neural Networks Bruges, pp. 467–472. Neurocomputing (2005)

Der, R.: Self-organized acquisition of situated behaviors. Theory in Biosciences 120, 179–187 (2001)

Der, R., Güttler, F., Ay, N.: Predictive information and emergent cooperativity in a chain of mobile robots. In: Artificial Life XI. MIT Press (2008)

Der, R., Liebscher, R.: True autonomy from self-organized adaptivity. In: Proc. of EPSRC/BBSRC Intl. Workshop on Biologically Inspired Robotics. HP Labs Bristol (2002)

Der, R., Martius, G.: The Playful Machine - Theoretical Foundation and Practical Realization of Self-Organizing Robots. Springer (2012)

Elbrechter, C., Haschke, R., Ritter, H.: Bi-manual robotic paper manipulation based on real-time marker tracking and physical modelling. In: IROS, pp. 1427–1432 (2011)

Friston, K.: The free-energy principle: a unified brain theory? Nature Reviews. Neuroscience 11(2), 127–138 (2010)

Friston, K., Adams, R.A., Perrinet, L., Breakspear, M.: Perceptions as Hypotheses: Saccades as Experiments. Frontiers in Psychology, 3 (2012)

Friston, K.J.: A free energy principle for biological systems. Entropy 14(11), 2100–2121 (2012)

Friston, K.J., Stephan, K.E.: Free-energy and the brain. Synthese 159(3), 417–458 (2007)

Grossekathofer, U., Barchunova, A., Haschke, R., Hermann, T., Franzius, M., Ritter, H.: Learning of object manipulation operations from continuous multimodal input. In: Humanoids, pp. 507–512 (2011)

Hauser, H., Ijspeert, A.J., Füchslin, R.M., Pfeifer, R., Maass, W.: Towards a theoretical foundation for morphological computation with compliant bodies. Biological Cybernetics 105(5-6), 355–370 (2011)

Hauser, H., Ijspeert, A.J., Füchslin, R.M., Pfeifer, R., Maass, W.: The role of feedback in morphological computation with compliant bodies. Biological Cybernetics 106(10), 595–613 (2012)

Jung, T., Polani, D., Stone, P.: Empowerment for continuous agent-environment systems. CoRR, abs/1201.6583 (2012)

Klyubin, A.S., Polani, D., Nehaniv, C.L.: Empowerment: a universal agent-centric measure of control. In: Congress on Evolutionary Computation, pp. 128–135 (2005)

Klyubin, A.S., Polani, D., Nehaniv, C.L.: Representations of space and time in the maximization of information flow in the perception-action loop. Neural Computation 19, 2387–2432 (2007)

Kolodziejski, C., Porr, B., Tamosiunaite, M., Wörgötter, F.: On the asymptotic equivalence between differential hebbian and temporal difference learning using a local third factor. In: Advances in Neural Information Processing Systems (2009)

Kolodziejski, C., Porr, B., Wörgötter, F.: Mathematical properties of neuronal td-rules and differential hebbian learning: a comparison. Biological Cybernetics (2008)

Kulvicius, T., Kolodziejski, C., Tamosiunaite, M., Porr, B., Wörgötter, F.: Behavioral analysis of differential hebbian learning in closed-loop systems. Biological Cybernetics (2010)

Lazar, A., Pipa, G., Triesch, J.: The combination of STDP and intrinsic plasticity yields complex dynamics in recurrent spiking networks. In: ESANN, pp. 647–652 (2006)

Lazar, A., Pipa, G., Triesch, J.: Emerging bayesian priors in a self-organizing recurrent network. In: Honkela, T. (ed.) ICANN 2011, Part II. LNCS, vol. 6792, pp. 127–134. Springer, Heidelberg (2011)

Markovic, D., Gros, C.: Self-Organized chaos through polyhomeostatic optimization. Physical Review Letters 105(6), 068702 (2010)

Markovic, D., Gros, C.: Intrinsic adaptation in autonomous recurrent neural networks. Neural Computation 24(2), 523–540 (2012)

Martius, G., Der, R., Ay, N.: Information driven self-organization of complex robotic behavior. PLOS ONE (in press, 2013)

Maycock, J., Dornbusch, D., Elbrechter, C., Haschke, R., Schack, T., Ritter, H.: Approaching manual intelligence. KI 24(4), 287–294 (2010)

Pfeifer, R.: "Morphological computation" - self-organization, embodiment, and biological inspiration. In: IJCCI (2012)

Pfeifer, R., Bongard, J.C.: How the Body Shapes the Way We Think: A New View of Intelligence. MIT Press, Cambridge (2006)

Pfeifer, R., Gomez, G.: Understanding Intelligence (1999)

Pfeifer, R.: Morphological computation – connecting brain, body, and environment. In: Sattar, A., Kang, B.-H. (eds.) AI 2006. LNCS (LNAI), vol. 4304, pp. 3–4. Springer, Heidelberg (2006)

Pfeifer, R., Lungarella, M., Iida, F.: Self-organization, embodiment, and biologically inspired robotics. Science 318, 1088–1093 (2007)

Prokopenko, M. (ed.): Foundations and Formalizations of Self-organization. Springer (2008)

Prokopenko, M.: Information and self-organization: A macroscopic approach to complex systems. Artificial Life 15(3), 377–383 (2009)

Prokopenko, M., Boschetti, F., Ryan, A.J.: An information-theoretic primer on complexity, self-organization, and emergence. Complexity 15(1), 11–28 (2009)

Ritter, H., Haschke, R., Steil, J.J.: Trying to grasp a sketch of a brain for grasping. In: Sendhoff, B., Körner, E., Sporns, O., Ritter, H., Doya, K. (eds.) Creating Brain-Like Intelligence. LNCS, vol. 5436, pp. 84–102. Springer, Heidelberg (2009)

Steffen, J., Elbrechter, C., Haschke, R., Ritter, H.J.: Bio-inspired motion strategies for a bimanual manipulation task. In: Humanoids, pp. 625–630 (2010)

Triesch, J.: A gradient rule for the plasticity of a neuron's intrinsic excitability. In: Duch, W., Kacprzyk, J., Oja, E., Zadrożny, S. (eds.) ICANN 2005. LNCS, vol. 3696, pp. 65–70. Springer, Heidelberg (2005)

Triesch, J.: Synergies between intrinsic and synaptic plasticity mechanisms. Neural Computation 19(4), 885–909 (2007)

Zahedi, K., Ay, N., Der, R.: Higher coordination with less control – A result of information maximization in the sensorimotor loop. Adaptive Behavior 18(3-4), 338–355 (2010)