Behavior as broken symmetry in embodied self-organizing robots

Ralf Der and Georg Martius

Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany
{ralfder—martius}@mis.mpg.de

Abstract

Self-organization—ubiquitous in nature—is a major challenge for both artificial life and modern robotics offering intriguing perspectives for practical applications utilized so far only incipiently. There is some progress, though, in formulating general objective functions for driving systems into self-organization (SO). Based on general principles like information maximization, these approaches are domain invariant and free of arbitrariness. However, and this seems to be a major source of concerns, 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 not justified. Instead, we 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. The approach is based on a new unsupervised learning rule staging two antagonistic activities—driving systems towards instability while preserving the physical symmetries of the system as much as possible. This leads to spontaneous symmetry breaking, the leading phenomenon of SO known from nature that has been overlooked by the robotics community so far. It is shown by a number of examples that the unsupervised learning rule induces an amazing variety of behaviors—patterns in space and time that can be interpreted as broken symmetries.

Introduction

Self-organization (SO) is a ubiquitous phenomenon in nature and a promising challenge to the creation of artificial autonomous systems. In particular, in embodied artificial intelligence, 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 about its bodily affordances and then, in a second step, one may focus on the exploitation of the emerging motion patterns—by guiding the SO process into the directions of potential benefits.

While there are many approaches toward structural SO, in particular self-assembly, the SO of behavior still is considered more as wishful thinking than as a true and systematic approach toward autonomy. A principled way toward the SO of behavior faces essentially two challenges. One is how to organize a robotic system in such a way that it starts to self-organize its behavior. Actually, the situation in that point is not too bad. There are several approaches based on formulating objective functions (OF) for SO. In recent years, several such OFs have been proposed, ranging from the maximization of predictive information Ay et al. (2008); Der et al. (2008); Ay et al. (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 and Stephan (2007); Friston (2012, 2010) 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) 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.

These OFs all fulfill the prerequisite of a principled approach to SO: as they are formulated in a domain invariant way, they do not determine specific directions for the autonomous development, avoiding to put in what one actually wants to get out. But this achievement creates a dilemma which is the second, more serious challenge to SO. In fact, 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 not justified. 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 this paper we study how a self-organizing approach to robot control can break symmetries of the robot-environment system such that structured behavior emerges. Our approach is based on a new learning rule, see Der (2013), applied to two robotic systems. By these examples, we want to make the reader aware of the phenomenon of spontaneous symmetry breaking that is in our opinion 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. It is one aim of this paper to contribute to the dissemination of this prospective ingredient for modern embodied robotics, see also Pfeifer and Bongard (2006); Pfeifer et al. (2007).

This paper is organized as follows: The next section describes the control framework which is the basis for the definition of the unsupervised learning rule in the following section. Then we present the first robotic scenario in section "Vehicles: behavior as broken symmetry" with examples of behaviors from sparse symmetry braking events. We formulate a "rule of thumb" for self-organizing behavior from symmetry breaking. It follows the section "The HEXAPOD" in which we study the emergent behavioral modes and control structures using a six-legged robot. Finally we conclude with a discussion. Supplementary material, especially videos, are available at `http://playfulmachines.com/ECAL2013`.

Control Framework

Fundamental to our approach is the closed loop control setup. The controller of the robot is given by a one layer feed-forward neural network transforming sensor values $x \in \mathbb{R}^n$ into motor values $y \in \mathbb{R}^m$ as

$$y = K\left(x, C, h\right) = g\left(Cx + h\right) \tag{1}$$

where C and h are the parameters (synaptic strengths and bias values, respectively) and $g_i(z) = \tanh(z_i)$ is the sigmoidal activation function. The translation between the external and the internal world can be done by a forward model predicting future sensor values on the basis of the current sensor and motor values. Here we use a linear network:

$$x_{t+1} = \phi\left(x_t, y_t,\right) + \xi_{t+1} = Ay_t + Sx_t + b + \xi_{t+1}$$

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 with the parameter matrices A and S, and the parameter vector b. The forward model can be adapted on-line by a supervised gradient procedure to minimize the prediction error as

$$\Delta A = \eta \xi y^\top \ , \quad \Delta S = \eta \xi x^\top \ , \quad \Delta b = \eta \xi \ . \tag{2}$$

In the applications, the learning rate η may not be small such that the low complexity of the model is compensated by a fast adaptation process. It is one message of this paper that, due to the strong interplay with the embodiment, these very simple control structures can produce amazingly complex behaviors.

One-Dimensional Example

Let us consider a wheeled robot on a rail with a single motor and a single wheel-counter measuring the wheel velocity. Connecting the simple controller given by equation (1) and interpreting the motor values as target velocities we can analyze the dynamical properties of the system. Let us first consider $h = 0$. For $C < 1$ there exists only one fixed point for $x = 0$, corresponding to the standing robot whereas for $C > 1$ there are two fixed points one for forward and the other one for backward driving. The system is fully symmetric in this respect assuming that also the morphology is perfectly forward-backward symmetric. More formally the system is invariant against inversion of the x-axis. For $h \neq 0$ there is an asymmetry in the bifurcation structure, which we will not discuss further, see Der and Martius (2012) for details.

At this simple example we can understand how symmetries can be broken by noise. Let the controller be given by $C = 0, h = 0$, such that the robot is in total rest. When we now increase C to a value larger than 1 we cross the bifurcation point and the resting state becomes unstable and the perturbations by e. g. noise decide to which fixed point the system goes.

Unsupervised learning for self-organization

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). In Martius et al. (2013), a modification of the PI, the so-called time-local predictive information (TiPI) was introduce for better coping with the problem of non-stationarity in continually learning systems. By maximizing the TiPI, a general learning rule for the synaptic strengths of a neural controller network was derived. Different from infomax principles derived so far, the method interrelates the principle formulated at the level of behaviors directly down to the synaptic level.

In Der (2013), starting from the learning rule given in Martius et al. (2013), a new rule was presented. Compared to the TiPI, this new rule was shown to drive the system toward self-organization in a more sensitive way, giving rise to a rich scenario of spontaneous symmetry breaking. This was argued to open ways to new classes of self-organized behavior. A discussion will be given below.

The rule is written as (all quantities are at time t)

$$\frac{1}{\varepsilon}\Delta C_{ij} = \delta y_i \delta x_j - \gamma_i y_i x_j \tag{3}$$

$$\frac{1}{\varepsilon}\Delta h_i = -\gamma_i y_i \tag{4}$$

where δx_t is the prediction error based on time $t-1$

$$\delta x_t = x_t - \phi\left(x_{t-1}, y_{t-1}\right)$$

or some other perturbation quantity[1]. δy_t is defined in terms

[1]The new rule is not restricted to using δx as the prediction error. Instead we are free to consider δx as any convenient change in or perturbation of the sensor dynamics. In the experiments described below we used the change of the sensor values in one time step.

of the world model as

$$\delta y_t = J^\top \delta x_{t+1} \tag{5}$$

where

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

is the Jacobian matrix of the forward model expressing the sensitivity of its output on the input y. In our linear model, we simply have $J = A$.

Moreover, γ_i is a neuron specific learning rate defined as

$$\gamma_i = 2\alpha \delta y_i \delta z_i \tag{6}$$

where α is an empirical quantity controlling the sensitivity with $\alpha > 1$, and $\delta z = C\delta x$ is the change in the postsynaptic potential caused by δx.

Discussion of the learning rule: self-induced symmetry breaking

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 term since it is the product of the input x_j into the synapse C_{ij} and the activation y_i of neuron i. As such it would strengthen all paths 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. This would drive the neurons into saturation. However, with the negative sign (and $\gamma_i > 0$, in standard situations), the term actually is anti-Hebbian, counteracting the saturation of the neurons.

The first contribution, $\delta y_i \delta x_j$ is Hebbian again, formulated, however, not in the activations itself but in their deviations from the predicted values as generated by the model. Given the relation between δy_t and δx_{t+1}, see equation (5), ΔC_{ij} is strengthened if there is a strong correlation between $\delta x_{t,j}$ and the components of δx_{t+1} being fed by $\delta y_{t,i}$. Roughly speaking, in that way the first term in the learning rule tries to increase the propagation of perturbations δx, driving the system towards instability. Here we can draw a parallel to homeokinetic learning (Der and Martius, 2012), where we also have two antagonistic terms which together should drive the system towards an exploratory behavior. The structure of the learning rules are similar but differ in details, as discussed below.

In the bifurcation scenario discussed above, the symmetry breaking was induced by changing the controller parameter manually. With the unsupervised learning rule, we have a self-referential system, a dynamical system that changes its parameters by itself, see also 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 drives the physical dynamics towards instability, eventually causing a **spontaneous** breaking of existing symmetries.

To give an example: consider a hexapod robot (see e. g. Figure 7 below) where the parameters C_{ij} represent the couplings between the sensors and the motors. Intuitively the δx and δy contain some information of the current mode of behavior, that is not already modeled perfectly by the forward model. Combined in the driving term the learning rule will amplify a latent (easy to excite) mode of behavior.

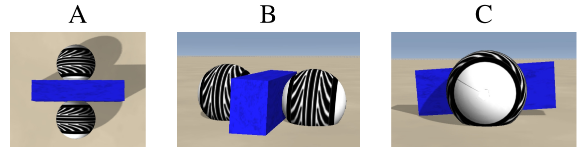

Figure 1: The TwoWheeled as a 3D physical object. The motor values are interpreted as wheel velocities, however the simulated motors have a maximal force. The ground is elastic so that the wheels sink into the ground depending on their load. Additionally there is friction and slip. The sensors measure the actual wheel rotation velocity. (A,B) wheel size 1: the robot is lying more or less flat on the ground when driving straight. 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 S1. Note the wheels are about 2% off center, so the robot is not fully forward/backward symmetric. (C) Wheel size of 1.2 the body can tilt to the front and back. These 3D effects would make both odometry and the execution of motion plans very difficult as they involve the full physics of the robot.

Vehicles: behavior as broken symmetry

Let us now apply the new learning rule to the specific example of a TwoWheeled robot (Figure 1) such that the characteristic properties of the self-organization process are illustrated. For the simulations the LpzRobots simulator (Martius et al., 2010) was used.

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 motors, 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 TwoWheeled case this means that all wheels are held at rest (velocity control). In robots with joint-position control, $y = 0$ means that all joints are driven towards their center position.

The combined system, comprising the physical and the synaptic dynamics, is fully deterministic. If starting in the least biased initialization the combined system may be in an unstable fixed point. We can either add small noise to the sensors for a short time interval or position the robot initially such that sufficient perturbation occur. Without further noise, the actual initial condition is fixed and the time evolution of the entire system is deterministic.

Symmetry breaking—a rule of thumb

Before going on 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 driving term in the learning rule ($\delta y_i \delta x_j$)) while conserving as much of the original symmetries of the system as possible. When only few of the symmetries are broken we call it the parsimony (or economy) 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 left-right and forward-backward symmetries. 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, walls, and ground). To give an example: If the robot drives in a straight line back and forth, the rotational symmetry of the space is broken, whereas the forward-backward, left-right symmetries are conserved. However, if the robot drives in a circle the rotational symmetry is conserved and the others are broken. So a 'good' behavior in the sense of parsimoniously broken symmetries would be driven in a circular pattern with both forward and backward driving.

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). 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 HEXAPOD further below.

Results

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 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. Remember that the learning rule gives no clue of how symmetries are to be broken.

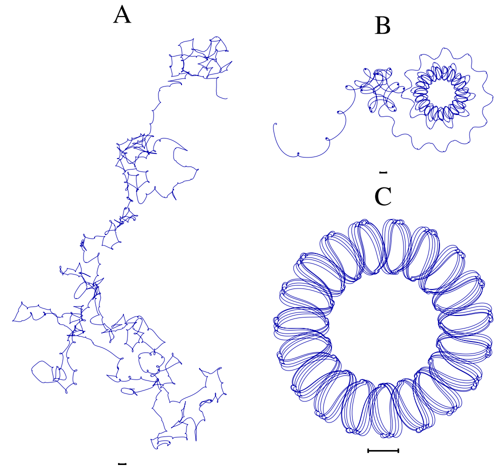

Figure 2: Deterministic trajectories of the robot in the ground plane emerging with different learning rates ε. If learning is fast ($\varepsilon > 0.01$), irregular trajectories occur (A). With lower rates (here $\varepsilon = 0.001$), after a transient phase of irregular motion through metastable attractors (B), the dynamics is converging toward a limit cycle behavior (C), called the master cycle below. The width of the robot is displayed by the small scale-line at the bottom. See also video S2. Parameters: $\alpha = 3$.

When using the controller (equation (1) with the learning dynamics given by equations (3) and (4) (and fixed forward model with $A = \mathbb{I}, S = 0, b = 0$, for simplicity), we expect the robot to start moving after some time[2] while trying to conserve as much of the original symmetries as possible. However, when using a learning rate ε above a certain value, 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 Figure 2(A). Still, note that these trajectories are fully deterministic. Nevertheless, our rule of thumb obviously is not valid in this regime as there is no visible footprint of the underlying symmetries—the invariances 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. Figure 2(B,C) is demonstrating a typical behavior of the robot. After starting, the robot is running through a kind of metastable patterns converging after some time toward a

[2] 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 HEXAPOD treated below, in the TWOWHEELED case, no substantial differences in the behaviors were observed.

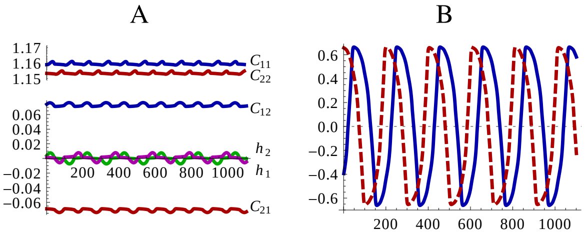

Figure 3: The circular pattern formation (Figure 2(C)) is hidden in the dynamics of the controller parameters as driven by the general learning rule. The C-matrix (A) is seen to be of a nearly perfect $SO(2)$ structure (rotation matrix), which can be described by a single rotation angle and a scaling, so instead of four parameters there are only two required. The h dynamics is seen to be periodic with a slight bias. (B) the two sensor values (wheel velocities).

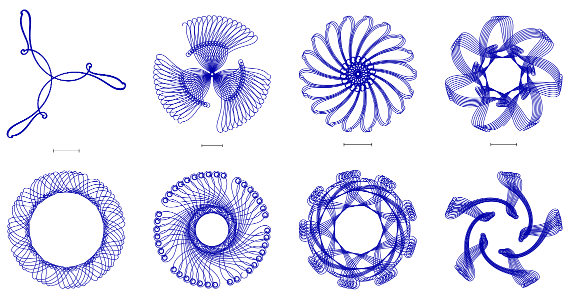

Figure 4: Patterns for frozen controller parameters occurring along the master cycle, Figure 2(C). Depicted is a selection of such patterns from parameter snapshots in one period of the state-parameter dynamics (about 200 time steps in Figure 3). If the learning is switched on again, the full dynamics is converging back to the master cycle.

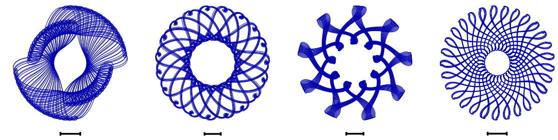

Figure 5: The emerging patterns also depend sensitively on the learning parameters. The figure shows the emerging patterns with changing α parameter (from left to right) as $\alpha = 1.0, 1.3, 1.9, 2.0$.

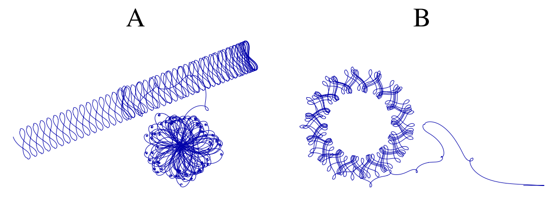

Figure 6: The role of embodiment. (A) Wheel size 1.1 (default 1.0). After a very irregular initial phase, the robot enters an aligned wiggly pattern, running at first to the right and then back toward the left lower corner. (B) Wheel size 1.125 leads to a circular pattern again. Parameter: $\varepsilon = .001$, $\alpha = 3$.

large scale circular pattern (CP).

The parameters of the controller during the CP (Figure 2(C)) are not constant but run themselves through a limit cycle as displayed in Figure 3. This is an immediate consequence of the close and persistent coupling between learning and physical dynamics. What happens if we keep them fixed at any time? The answer is quite astonishing: a variety of different patterns emerges, as displayed in Figure 4. This also illustrates that the parameter dynamics within the limit cycle is actually important for the particular pattern. The former can be seen as a transient along the many stable patterns with fixed parameters. Upon switching on learning again, the system rapidly returns to the original CP (with a different spatial position). This so called pattern spin-off effect was for the first time reported in Der (2013), this paper presents additional results demonstrating the richness of that phenomenon.

At present we do not have a complete microscopic understanding of the effects. Still, at the level of phenomena, there is a number of observations. One is that the very nature of the emerging patterns depends in a most sensitive and intricate way on both the embodiment and the learning dynamics. For instance, by varying the so called sensitivity parameter α (equation (6)) of the learning rule we obtain a set of quite different CPs as shown in Figure 5.

Alternatively, we may change the embodiment and obtain another class of behaviors. One option is to increase the wheel size that causes the trunk of the robot to tilt more when accelerating, see Figure 1. Two exemplifying trajectories are presented in Figure 6. For certain wheel sizes we may get also linear patterns, as they are predominant with a fully forward/backward symmetric robot. On the general level we may argue that for the linear pattern not the rotation symmetry is partially conserved, but the translational one along the line. However only a small change in the wheel size yields a CP again but with a very different fine structure.

Are we lost? Confronted with such an overwhelming variety of emerging patterns, are we faced with a robot that is behaving completely unpredictable confirming just the concerns against self-organizing robots we wanted to dispel? One answer is found by taking a look at the controller parameters. As Figure 3 shows, the controller matrix C is of a very specific structure, it is a nearly perfect (scaled) rotation matrix. Any such matrix rotates a vector by an angle and stretches it by a factor, so it is parametrized by only two

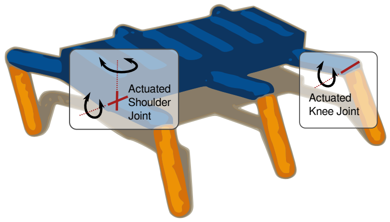

Figure 7: The HEXAPOD. 18 actuated DoF: 2 shoulder and 1 knee joint per leg. Fully forward/backward and lateral symmetric.

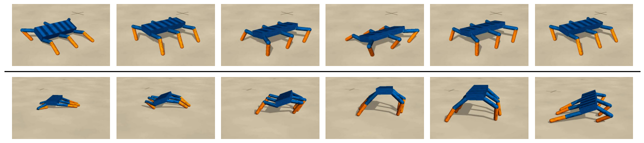

Figure 8: Initially, after about 20 min the robot develops a swaying motion pattern (top row), as if it is very actively trying to move the legs in a coherent way while keeping ground contact. 50 min later a raising behavior develops where the trunk is repeatedly being lifted from the ground.

variables. This specific structure of the controller matrix, obtained by the learning, also seems to be responsible for the specific pattern creation effects. In principle, there are essentially three possibilities for the asymptotic system dynamics with a fixed controller matrix: fixed point, limit cycle and chaotic attractor. As we have observed in a series of experiments, the limit cycles are most likely to occur with a rotation matrix. This is intuitively understandable given that the physical space is invariant against rotations.

In a sense, this specific controller structure is like a footprint left by the symmetries of physical space, imprinted into the controller by **spontaneous** symmetry breaking, driven by the unsupervised learning procedure that does not break any symmetries by itself.

So, from looking inside, there is a coarse explanation of how the robot achieves the patterns—by discovering, so to say, the world of rotation matrices. However, the point of major interest is that the learning finds those specific structures. On a general level, an understanding may be given by the rule of thumb: a pattern in space can only emerge from breaking the spatial symmetries inherent in the physics of the robot. 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 actual patterns emerging in the learning scenario are not circles but CPs. Nevertheless, they are still invariant against rotations about a definite angle, see in particular the patterns of Figure 4 and Figure 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.

The HEXAPOD

Let us now follow the trace of symmetry breaking with a high-dimensional six-legged robot: the HEXAPOD, see Figure 7. We choose this robot because it will be seen to reveal symmetry breaking phenomena in a particularly clear way. The robot has six legs, each one with three degrees of freedom (DoF). Each of the 18 joints is actuated by a servo motor and contains a sensor that is measuring the actual joint angle. The effective torques acting on the joint axes are determined by a PID controller with a limited force. To enable a body feeling (some useful feedback from the interaction), this force limit is proportional to the deviation from the set point, such that there is an elastic reaction to external forces, similar to 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 the motor values are zero ($y = 0$) 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 δx that is further amplified by the learning dynamics.

What can we expect to happen? Depending on the concrete situation (e. g. particular ε) different behaviors may emerge. In most cases the robot starts with a swaying motion pattern, see Figure 8 and video S3. 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 physical constraints enforced by the ground contacts.

More interesting behaviors emerge after some time, for instance a raising behavior, see Figure 8. The entire development can be followed in short pieces in the videos S3-S6. There is another surprise—when looking at the parameters of the controller. In the TWOWHEELED case the C-matrix developed into a rotation matrix. 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 interaction with the ground.

Yet, as Figure 9(A) shows, the emerging sensor-to-motor coupling matrix is highly structured, reflecting the original symmetries to a high degree. Both the shoulders vertical direction and the knees are seen to follow essentially the same strategy for moving the body. This is in agreement with our rule of thumb since this collective strategy allows the body to be moving, but with a maximum degree of coherence between the individual constituents of the body. Moreover, the coupling matrix reveals the whole-body nature of the behavior—the control for each body part is generated by

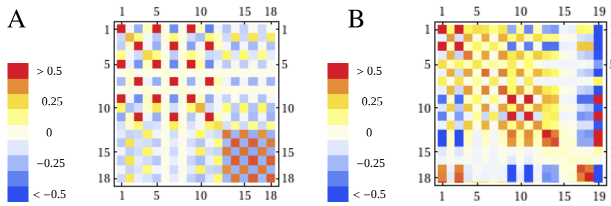

Figure 9: The parameters of the controller (C-matrix) for two scenarios. Element C_{ij} represents the coupling from sensor j to motor i. Indices: 1-12: shoulders (vertical/horizontal), 13-18: knees. (A) Swaying motion, Figure 8. (B) Seesaw motion with velocity sensor (index 19), see Figure 10. The difference between the swaying and seesaw behavior are clearly visible in the structure of the matrices. While in swaying all legs follow the same strategy, the antiphase nature of the seesaw behavior is reflected in the different sign distribution of the matrix elements.

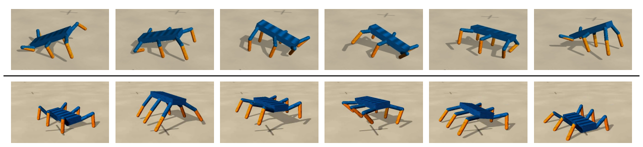

Figure 10: Seesaw motion pattern with forward/backward speed sensor (top row), see video S7. Jumping motion pattern emerging with vertical speed sensor (bottom row), see video S8. Note the robot is in the air in the frames 2-4.

combining both excitatory and inhibitory signals from the sensors of all joints in a systematic manner.

The formative power of exteroceptive sensors

Up to now we were using only proprioceptive sensors so that the orientation of the robot in space can only be measured very indirectly, e. g. by additional load to the joint motors due to gravity forces. By including exteroceptive sensors, the development of the modes can be influenced and driven into desired directions. Adding a sensor measuring the forward velocity of the robot generates a seesaw motion. In contrast, a sensor measuring the vertical velocity of the robot leads to a pronounced jumping behavior, see Figure 10. Also here, we find highly structured controller matrices, see Figure 9(B) for the seesaw case. Note the strong coupling of the exteroceptive sensor to the motor neurons showing the functional role of that sensor. It distinguishes forward and backward motion and thus this symmetry is lost, so that in the learning process behaviors with broken forward-backward symmetry are favored.

Discussion

This paper tries to answer essentially two questions. The first question is about how to organize self-organization, in other words, 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 equations (3) and (4), which fulfills the main criterion for a genuine SO: it is universal in the sense that the only necessary information about the system is given by the number of sensors and that of motor neurons, any further information being acquired by the co-learning self-model in a bootstrapping process.

The second question we want to answer in this paper 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 which extent 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 which induces, given the constraints, corresponding symmetries in the physical system. 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 steps, following—similar to what we know from nature—a kind of parsimony principle. This leads 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 acts as a pattern factory: the parameters occurring along the limit cycle produce a great variety of spin-off patterns. While this effect has already been reported in Der (2013), this paper presents further results and gives additional insights into this interesting effect.

Continuing the work started in Der (2013), similar effects of symmetry breaking were obtained in the example of the Hexapod. We observed the excitation of body related, high activity modes with a high degree of coherence between the body parts. These modes were argued to be in nice agreement with 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 argument, it was often not clear what the robot is actually doing. With the new learning rule and the concept of behaviors as broken symmetries, this is now (a little) different. The essential difference between homeokinetic learning and the ULR is the dynamics with the least biased initialization ("do nothing" region with all synapses zero). While the time-loop error of homeokinesis has a pole there, the infomax based objectives are smooth in that region. It is basically this smoothness that makes the learning to integrate the responses of the system dynamics in a sensitive way. As compared to the TiPI (Martius et al., 2013), the learning dynamics used here is even smoother and even more concentrated on system responses which explains the prevalence of spontaneous symmetry breaking effects. At a more formal level, we see the difference also in the drift of the local Lyapunov exponents: while homeokinesis drives small exponents stronger than larger ones, the situation is inverted in the present learning dynamics. Given the formative interplay between state and learning dynamics, this has important consequences for the emergence scenario of the behaviors.

The principles and examples given in this paper—in particular the emergence of coherent modes, the TWO-WHEELED as a pattern factory and the various modes realized by the HEXAPOD—may help us to better understand and exploit the synergy between embodiment and SO of autonomous robots.

Acknowledgments

The authors gratefully acknowledge the great hospitality in the group of Nihat Ay. We thank for many helpful and clarifying discussions with Nihat Ay, Nils Bertschinger, Eckehard Olbrich, and Keyan Zahedi. GM was supported by a grant of the DFG (SPP 1527).

References

Anthony, T., Polani, D., and Nehaniv, C. L. (2009). Impoverished empowerment: 'meaningful' action sequence generation through bandwidth limitation. In *ECAL*, pages 294–301.

Ay, N., Bernigau, H., Der, R., and Prokopenko, M. (2012). Information driven self-organization: The dynamical systems approach to autonomous robot behavior. *Theory Biosci.*, 131(3):161–179.

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

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

Der, R. (2013). On the role of embodiment for self-organizing robots: behavior as broken symmetry. In Prokopenko, M., editor, *Guided Self-Organization: Inception*. Springer, to appear.

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

Der, R. and Liebscher, R. (2002). True autonomy from self-organized adaptivity. In *Proc. Workshop Biologically Inspired Robotics*, Bristol.

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

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

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

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

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

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

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

Martius, G., Der, R., and Ay, N. (2013). Information driven self-organization of complex robotic behaviors. preprint, submitted to PLoS ONE.

Martius, G., Hesse, F., Güttler, F., and Der, R. (2010). LPZROBOTS: A free and powerful robot simulator. robot.informatik.uni-leipzig.de/software.

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

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

Prokopenko, M., editor (2008). *Foundations and Formalizations of Self-organization*. Springer.

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

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