Adaptive **Behavior**

Attractor Landscapes and Active Tracking: The Neurodynamics of Embodied Action

Mario Negrello, Frank Pasemann
Fraunhofer Institute for Intelligent Analysis and Information Systems, Sankt Augustin, Germany

Behavior is the product of three intertwining dynamics: of the world, of the body and of internal control structures. Neurodynamics focuses on the dynamics of neural control, while observing interfaces with the world and the body. From this perspective, we present a dynamical analysis of embodied recurrent neural networks evolved to control a cybernetic device that solves a problem in active tracking. For competent action selection, agents must rely on the attractor landscapes of the evolved networks. Insights into how the networks achieve this are given in terms of the network's dynamical substrate, which highlights the role of the network's inherent attractors as they change as a function of the input parameters (sensors). We introduce some terminological extensions to neurodynamics to allow for a more precise formulation of how attractor changes influence behavior generation: in particular, attractor landscapes, which are the space of all attractors accessible through coherent parametrizations of the network (input stimuli), and the meta-transient, which resolves behavior by approaching attractors as they shape-shift. We apply these concepts to the analysis of interesting behaviors of the tracking device, such as temporal contextual dependency, chaotic transitory regimes in moments of ambiguity, and implicit mapping of environmental asymmetricities in the response of the device. Finally, we discuss the relevance of the concepts introduced in terms of autonomy, learning, and modularity.

Keywords neurodynamics · dynamical systems · active tracking · chaos · recurrent neural networks · cognition

1 Introduction

1.1 Attractor Landscapes and Behavior

How many behaviors does an agent have in stock? A starting point for assessing this question might be the observation that behavior is the product of three intertwining dynamics: of the body, of the environment, and of internal control structures (Varela, 1979). Clearly, the complexity of behavior can originate in any one of these sources or in all three; as Simon (1969) pointed out, the complex path an ant takes might be a function of the surface it walks on, the ant's body, or its neurons. The complexities of behavior resulting from the environment and body physics are easier to observe and explain, but the complexities ensuing from neural control remain, to a large extent, mysterious.

So, let us consider that control structures are responsible for behavior, immersed between the dynamics of the body and of the environment. Neurodynamics (Negrello & Pasemann, 2008; Pasemann, Steinmetz, Hülse, & Lara, 2001) as a theory provides tools,

Correspondence to: Mario Negrello, Fraunhofer Institute for Intelligent Analysis and Information Systems, Department of Autonomous Robots, Schloss Birlinghoven, 53754 Sankt Augustin, Germany.
E-mail: mario.negrello@iais.fraunhofer.de; mjn67@cornell.edu
Figures 1–5, 7–10, A1 appear in color online: http://adb.sagepub.com

(2008), Vol 16(2/3): 196–216.
DOI: 10.1177/1059712308090200

extended from dynamical systems theory (DST), for the analysis of dynamical entities resulting from embodied control structures. Discrete time recurrent neural networks (RNNs) implement control of the embodied agents studied here. When parametrized[1] by the input, RNNs are in fact a collection of dynamical systems, in which a large repository of dynamics might be encountered. We call the dynamical repository given by one RNN and its possible parametrizations the "attractor landscape."

We argue that meaningful behavior results from the exploitation of multifaceted dynamical entities by transients. In DST, transients are defined as the sequences of states of a dynamical system in phase space before asymptotic states (attractors) are reached. However, during behavior it is unlikely that transients are given time to settle on one attractor, because of ever-present feedback, both from the body (as in proprioception) and the environment (exteroception, stimuli). Rather, behavior hops under their influences and, regarding the input to the system as parameter changes, both in the form of feedback or perturbations, results in the realization of a class of dynamical systems, upon which parametrizations the attractor landscape is dependent. Therefore, we find it informative to think of behavior as happening on a "meta-transient." In contrast to a transient, which approaches an attractor in one dynamical system, a meta-transient is, roughly, the transient that is subject to varying parametrizations (changing stimuli) and therefore to distinct dynamical systems. In other words, the meta-transient explores the basins of attraction as they change (in size, number or shape) under the influence of the changing stimuli, internal state, and situations. In fact, our thesis is that behavior can be regarded as being generated by a meta-transient across the attractor landscape of a parametrized dynamical system (an RNN) that acts non-trivially through parameter shifts.

The multifarious characteristic of behavior is then a consequence of the changes in attractor structures of the activity space that influence the meta-transient. Because the portions of the attractor landscape accessible by the network change as a function of the input, so does the overlying meta-transient. Moreover, although in principle attractors fit into a small number of categories (fixed-points, cyclic attractors, quasi-periodic attractors, and chaotic attractors), their characteristics are far from exhausted by this categorization. In some sense, the shape of an attractor is its identity.[2] Therefore, an RNN might contain a potentially immense number of different attractors (Berry & Quoy, 2006). Regarding the attractor landscape as a result of input, attractors approached by transients might either change smoothly (which we later refer to as attractor morphs) or "catastrophically," depending on whether a bifurcation boundary is crossed. Here, we also introduce the notion of "paths in parameter space." Intuitively, these are sequences of stimuli that are correlated through embodiment and the environment, which is structured and coherent. The implications of this terminological refinement are exemplified by the results to come.

In an embodied agent, the transients are driven by the input to exploit the dynamical substrate of an RNN as a scaffold for behavior expression (Beer & Gallagher, 1992). In what follows, we demonstrate empirically how this exploration can occur, using a toy problem in active tracking. The goal of the depictions included is to highlight the dynamical entities as reductionistic concepts for the explanation of behavior. However, before the introduction of the problem itself, we review some necessary neurodynamics terminology, while introducing some new terminology.

1.2 Definitions

We look briefly at the definitions of the terms that appear most often in the article. For more detailed and comprehensive formalizations of parametrized dynamical systems and related concepts, we refer to Pasemann (2002).

By assumption, we (and others) consider input to the network to be equivalent to parametrizations, which change slowly in comparison with the dynamics of the network. In an embodied problem, these parametrizations carry the structure of the interactions. So, following Varela (1979), we write "structural coupling" ($\rightleftharpoons$) of the agent ($\mathfrak{a}$) and environment ($\mathfrak{E}$) at a given moment (t) as leading to ($\rightsquigarrow$) a set of parameters (ρ) to the control structure, the RNN:

$$\langle \mathfrak{E} \rightleftharpoons \mathfrak{a} \rangle_t \rightsquigarrow \rho_t. \tag{1}$$

A "path in parameter space" $\mathfrak{P} = (\ldots, \rho_{t-1}, \rho, \rho_{t+1}, \ldots)$ refers to the temporal sequence of parameter changes resulting from the structural coupling with the environment, where the reticences before and after the parameters stress the ongoingness of behavior.

In a parametrized RNN, such as those we handle here, every parameter defines a distinct "dynamical system" $\mathfrak{D}_\rho$. We note that a discrete time dynamical system is a map that takes a state in phase space to the next. So, one parameter (ρ) given by the interaction of $\mathfrak{a}$ and $\mathfrak{E}$ will take state s_t to s_{t+1}:

$$\mathfrak{D}_{\rho_t}(s_t) \mapsto s_{t+1}. \tag{2}$$

Any state s of the network's activity in phase space might, or might not, have reached an "attractor." An attractor of a dynamical system is an asymptotic state. A dynamical system might possess "coexisting attractors," in which initial conditions in different basins of attraction will lead to distinct attractors. The denomination of the state of a dynamical system before the attractor is reached is the "transient state."

It is worth noting that when the network has motor efference, a transient may also be a functional state, which produces a change in the state of $\mathfrak{a}$, possibly an action, changing the relations between $\mathfrak{a}$ and $\mathfrak{E}$, as in Equation 1, recursively resolving subsequent parametrizations. However, it is not necessary that the network should only produce behavior if an attractor has been reached. If the state of the network is constantly projecting to motor output, a transient state might also produce an action. Therefore, the path in parameter space $\mathfrak{P}$ is determined by the sequence of interactions between the agent and the environment, which, although influenced by the attractors, is not strictly a function thereof.

Consequently, the "meta-transient" $\mathfrak{M}$ is simply the set of states of the network in phase space as it is constrained by the structural coupling and its respective parametrizations, $\mathfrak{M} = (\ldots, s_t, \ldots)$. However, this can be rewritten so as to outline the fact that the states are dependent on the application of the map $\mathfrak{D}_{\rho_t}$ with the parametrization that changes over time:

$$\mathfrak{M} = [\ldots, \mathfrak{D}_{\rho_{t-1}}(s_{t-1}), \ldots]. \tag{3}$$

Note that a "bifurcation" is a region of structural instability, where the dynamics of two neighboring dynamical systems have qualitatively different attractor sets. Paths in parameter space $\mathfrak{P}$ can be within or across the bifurcation boundaries. In the first case, attractors undergo smooth changes or "morphing," while bifurcation boundary crossings lead to qualitatively distinct dynamical behavior.[3] To illustrate this, in Figure 1 we have plotted a bifurcation sequence respective to the change of input to neuron (θ_2) of an example two neuron network.[4]

Above this bifurcation diagram, we show a plot of the Lyapunov exponent respective to it, which, when positive, indicates chaotic attractors.

The dynamical substrate mentioned in the Section 1 is the attractor landscape of an RNN, which is determined by its structure (i.e., the weights and biases of one RNN and its possible input parametrizations, which, as we have seen, are constrained by the structural coupling between an agent and its environment). We assume that the structural variables of the network (weights and biases) remain fixed for the duration of a trial. For instance, the bifurcation sequence of Figure 1 is the projection of the attractor landscape along the dimension of θ_2.

Consequently, for one embodied network with a fixed structure, the agent will do that which the network structure plus parametrizations allows; the agent will act according to its attractor landscape subject to paths in parameter space. The attractor landscape is, so to speak, the behavior invariant of the embodied agent. In other words, the capacity for behavior is given by the attractor landscape and the interaction with the environment. In the following sections, we aim to substantiate this claim.

1.3 Toy Problem in Active Tracking

To illustrate the relevance of the concepts mentioned above, we present a toy problem in active tracking, in which evolutionary robotics (Harvey, Di Paolo, Wood, Quinn, & Tuci, 2005) methods beget the structural parameters of the networks (weights and biases). Simply put, the problem is that a cybernetic head should be able to follow, with its gaze, a ball that bounces irregularly in a two-dimensional plane (i.e., within a frame; see Figure 2). The primary problems of the head are to know in what direction to turn and with what velocity. A secondary problem is to actively search for the ball, in case it is lost from sight.

Our problem then is to discover the dynamical entities that allow the network to do this. That is, the focus of the analysis is to see how the meta-transient might hop between attractors while attractors change as a function of changing input patterns, and how

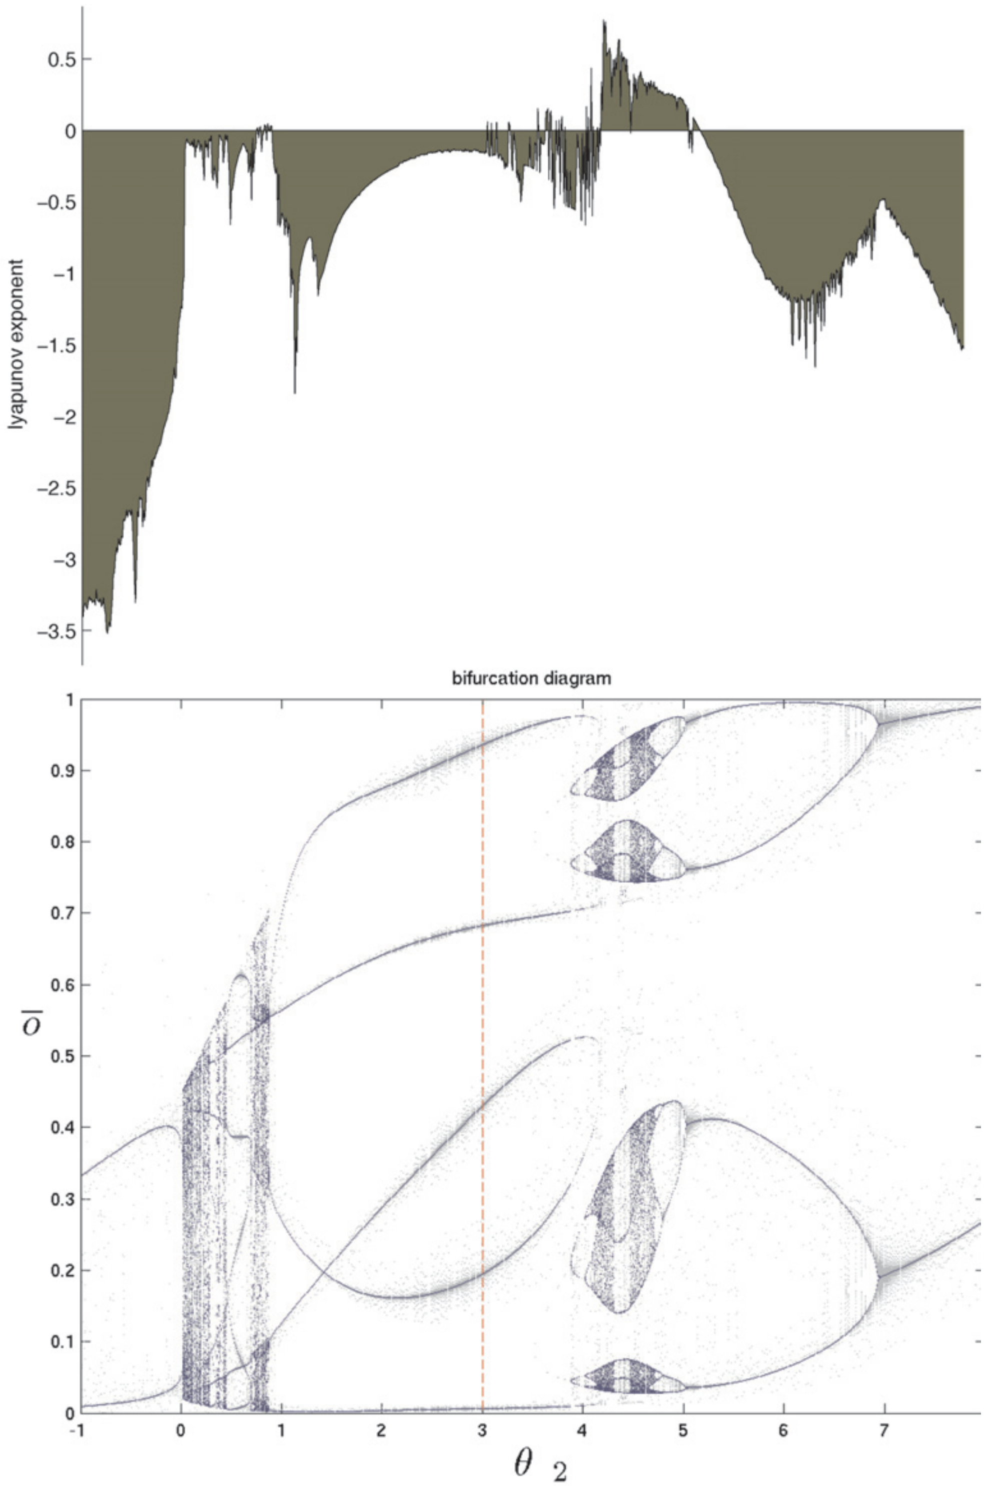

Figure 1 Bottom: Bifurcation sequence for a two-neuron network parametrized by input to unit 2 (θ_2). We plot the mean output activity of the two neurons. As the parameter is varied, the network undergoes qualitative changes in dynamical regimes. The scattered dots are transient states (first 300 iterations), and the thick lines belong to the attractors (next 100 iterations). The region around the dashed line has a period 5 periodic attractor. Close to bifurcations, the smearing of gray dots indicates longer transients. Top: The associated Lyapunov exponents where positive values indicate chaos. Around θ_2 = 4, a chaotic attractor coexists with a period 5 attractor.

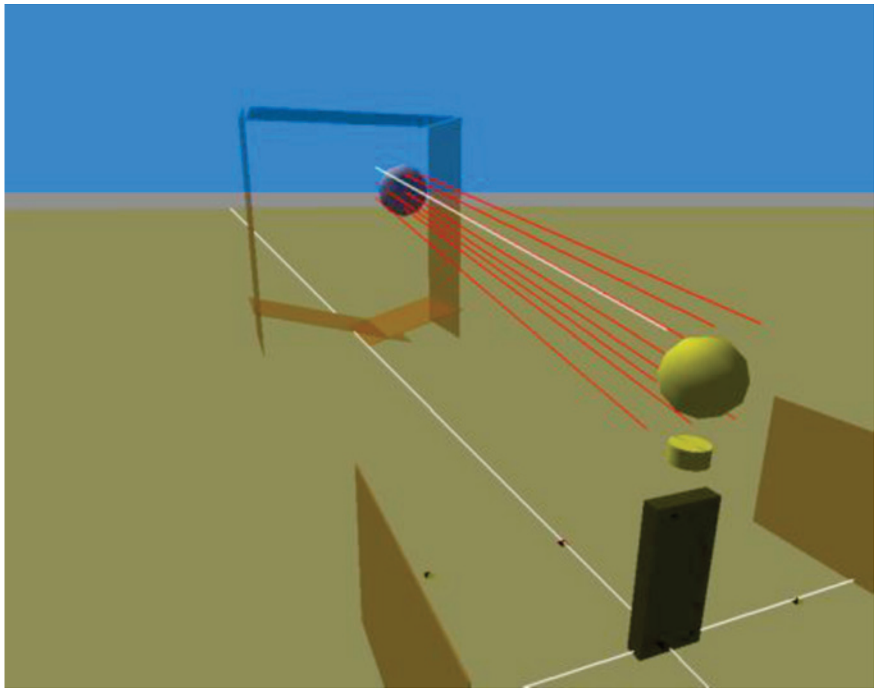

Figure 2 The simulated environment and the agent. The lines represent the distance sensors, where white means that the ray does not encounter the ball. The head is mounted on a neck controlled by two motors for pitch and yaw (the distance from the head to the ball is 3 m).

the projection of attractors becomes proper motor action.

The main motivation for this formulation is to have a highly dynamical problem where embodiment by necessity plays a fundamental role. In the case under study, inertial factors are crucial components of behavior, which impact on the dynamics of the sensory motor loop. In this way, we can observe the behavior that occurs between attractors, in the meta-transient.

In Section 3, we describe in more detail the useful characteristics of problem solving employed by the evolved networks, and the respective dynamical entities responsible. Here, we outline the dynamical features of the tracker's problem solving that are addressed in Section 3.

1. The projections of the high-dimensional attractors onto the motor units are at the core of the solution. We speak of projection shapes as the action identity and indicate how morphing may also be responsible for proper action selection.
2. Control is established by negative feedback, where the control signal is provided by the environment; the action is supported by the attractor landscape, which determines the amplitude of the oscillations.
3. The agents display temporal context dependency. We show that there may be two coexisting attractors associated with one parametrization of the network, and that these attractors may lead to different actions; a path in parameter space usually keeps the meta-transient under the influence of the same basins of attraction.
4. Robustness to errors and mismatches appears because of multiple possible network responses that generate similar actions to a given pattern. This means that the attractor landscapes also have redundancy in action. Conversely, there might be different actions to one stimulus pattern in ambiguous circumstances. This may also be due to chaotic attractors in ambiguous regions of the attractor landscape.
5. Using an analysis of the meta-transient, it becomes possible to infer an implicit modeling of the physical characteristics of the environment, such as gravity.

Figure 3 A sequence of 15 frames of stimuli, gathered during an average trial. The stimuli are analogous to a nine-pixel retina.

2 Methods

2.1 Description of the Problem

The toy problem here consists of a cybernetic tracking device following a ball in a virtual environment (see image of the environment and agent in Figure 2). The task of the device (the head) is to track a ball that bounces within a frame, keeping the ball in sight at all times. The device is a head mounted on a neck, composed of two motors for yaw and pitch. The input sensors are an array of nine distance sensors that signal linearly in the range 0 (nothing meets the ray) to 0.15 (ball closest to the head). The head is directed towards a frame that constrains a ball, which hops erratically in the two-dimensional plane. The ball is subject to gravity and to the geometry of the frame. As mentioned, by design the ball does not lose energy as it bounces. The head "sees" the ball as the interaction between the distance rays and the ball. It is possible to get an idea of the type of input to the head by looking at Figure 3, where 15 sequential inputs gathered during a trial are depicted (the input is analogous to a nine-pixel retina). The simulated neck motors take the desired velocity as input, up to a certain force [according to the specifications of the simulation library of rigid body dynamics used in the simulation, the famed open dynamics engine (ODE); see Table 1 for other physical properties of the simulation].

As is usual in evolutionary robotics, an evolutionary algorithm selects the RNNs according to a fitness function defining the aptitude of the agent to solve the prescribed problem. So, the objective of the head (keeping the ball in sight) is reflected in the fitness function: the more the agent is able to keep the ball in sight (the sensor input accumulated over the trial time) while minimizing oscillations, the fitter the network (see Equation 4).

The fitness function of a trial for an individual ("ind") is

$$fitness_{\text{ind}} = \underbrace{\alpha \sum_{t=1}^{c} \sum_{s=1}^{9} S_s(t)}_{LT} \underbrace{- \beta \sum_{t=2}^{c} \sum_{m=1}^{2} [M_m(t) - M_m(t-1)]^2}_{QT} . \quad (4)$$

Here, $S_s(t)$ is the input at sensor neuron s at time t and $M_m(t)$ is the activation of the motor neuron m at t. The first term (LT) denotes the sum of stimulus input per

Table 1 ODE simulation physics.

Entity	Property	Quantity
Head	Mass	3 kg
Head	Height	2 m
Yaw and pitch motor	Maximum force	5 N
Yaw and pitch motor	Maximum velocity	90° s^{-1}
3 × 3 distance input array	V and H distance between sensors	0.25 m
Ball	Radius	0.5 m
RNN and physics	Update frequency	100 Hz

cycle, while the negative quadratic term (*QT*) aims to minimize oscillations. α and β are parameters to weight the terms and can be altered on-the-fly during evolution, depending on the experimenter's emphasis. c denotes the number of steps or cycles in a single trial. Note that all the knowledge the fitness function requires is that about the motor activities, so all the knowledge of the fitness function is also available to the agent.

The model used in the experiments is a discrete time RNN, with the hyperbolic tangent as the nonlinear transfer function. The input layer is composed of linear buffers and receives no backward connection. In Appendix A we give the weights of two of the evolved networks as examples.

$$a_i(t+1) = \sum_j^n w_{ij}\tau[a_j(t)] + \theta_s;$$

$$i, j = 10, \ldots, n; \quad s = 1, \ldots, 9. \quad (5)$$

Here, a_i is the activity of the ith unit of the network, the total number of units is n, and τ is a sigmoid function. In this case, the hyperbolic tangent, w_{ij}, reads i receives from j with weight ditto. θ_s is interpreted as a slowly varying input from the sensors. The sensors are units 1–9, and the motors are 10 (yaw) and 11 (pitch).

2.2 Challenges for the Tracker

Although a toy problem, the solution is not trivial, because of the dynamical physical simulation. Figure 2 shows a depiction of the environment, and the difficulties arising from it are listed, as follows.

1. The head has to cope with rather meager input (a mere nine distance sensors) and, moreover, every pattern taken individually is ambiguous (is the ball coming into or escaping from view?). Also, even small changes in ball position relative to the sensor rays might lead to large input changes (e.g., in one cycle the input of one sensor might drop from positive to zero, as can be seen in Figure 3).
2. The head's foe, the bouncing ball, is designed to bounce erratically as a result of being dropped from different initial positions and because of the different angles of the bottom platforms (see Figure 2). For most initial conditions, the bouncing trajectories of the ball are highly unpredictable.
3. By design, the ball does not lose energy as it bounces, implying that when the ball bounces in different positions of the side walls or bottom platforms, it has very different velocities in the z and y Cartesian axes. For example, when the ball bounces sideways in the bottom of the frame, the horizontal velocity is much higher than when it bounces higher up in the frame. The ball is constantly subject to gravity of 9.8 m s^{-2}.
4. The network has no knowledge of the frame, so in principle the head has no information about the exit angle of the ball after it bounces against the frame. This means that if the network has a stereotypical response to ball following (such as pure asymmetry of the left–right weights), it is bound to lose track, as observed in the first generations (approximately up to 50 generations).

2.3 Evolution of Networks

For the artificial evolution, we use the ENS^3 algorithm (Dieckmann, 1995). The genome of the evolution is the structure of the networks themselves. The variation operator adds or deletes units (neurons) and synapses, as well as changing weights according to on-the-fly specified probabilities. The experimenter can also limit the number of units in the network, as well as introducing costs for extra units. The selection of the agent that generates offspring is rank-based. The offspring production is a Poisson process controlled by the shape of a distribution. This is done in order to keep a high diversity of the networks.

All parameters of the evolution are alterable on-the-fly (i.e., during the evolution itself; Huelse, Wischmann, & Pasemann, 2004). In this way, the experimenter is enabled as a "meta-fitness," selecting agents not only by their raw fitness but also by remarkable characteristics of behavior. The controllable parameters are the weights of the terms of the fitness function, the shape of the distribution of offspring production, the number of cycles of the trial, the number of initial conditions (such as the initial position of the ball), the number of neurons and synapses, and so forth. Off-line parameters are, for example, the frequency of the virtual simulation and the refresh rate of the networks.

2.4 Motor Projections of Attractors

Effective methods of analyzing small neural networks for their dynamic capacity have shown how paths in parameter space are able to change the behavior of the network both qualitatively and quantitatively. The presence of bifurcation boundaries in parameter space illustrates how even slight changes of a parameter can bring the networks to different dynamical regimes. Unfortunately, the powerful mathematical methods that can be used with very small networks become very complicated with highly recurrent networks, such as this one. Nevertheless, in an ambitious project recently, a systematic analysis of the dynamics of small continuous time RNNs has been advanced (Beer, 2006).

Nonetheless, complex high-dimensional spaces are not wholly intractable. Therefore, we rely on projection methods for the analysis, such as projecting the many dimensional orbits to motor space according to the definition of connected paths in parameter space, which indicate bifurcations as well as attractor morphs. Motor projections point to many of the relevant aspects of the dynamical entities involved, and provide for intuitive visualization of the meanings of the meta-transients. Figure 8, for instance, depicts a considerable part of the action set to the whole set of possible inputs. Although this analysis sacrifices, for example, the determination of precise bifurcation boundaries, it nevertheless allows for a bird-eye's view of the totality of the agent's action set.

Problem-specific knowledge can also simplify the input dimensions. In our case, we reduce the nine-dimensional input space to the two principal components of the stimulus, which are the relative positions of the ball to the sensory array in the horizontal and vertical dimensions. This reduction has the added benefit of showing how raw numbers of sensory input can drastically be reduced by the mild assumption that the changes in input space are correlated. This is because both the embodied agent and the environment are extended in the world, and so the inputs are never scrambled. The possible configurations of the sensory stimulus in our case are constrained to the possible interactions with the rays and the ball, which define the sensory manifold, and thereby all the possible paths on the sensory space (as in Philipona, O'Regan, Nadal, & Coenen, 2003). Through these defensible simplifications, it became possible to analyze the "action set" of the agent, despite the high dimensionality of both the input and the internal states.

3 Results

3.1 Tracking Behavior Across Attractors

Competent behavior occurs when the head is able to match its velocity with the velocity of the ball. This requires modulation of both the direction and the force applied to the motors. From the input alone, neither direction nor force are decidable. So, in order to have the best tracking, the past states also have to be taken into account for the current action. This means that for the optimal solution the system will also require memory, found in the internal states of the hidden layer. As we see, it is the profile of the motor output wave that modulates both the force and direction of the neck. The choice of wave profile (the attractor translated into a motor output time series; see Figure 5) for control is equivalent to the choice of an attractor invoked by the sensory-motor loop.

3.2 Solutions

The first solutions (before the 50th generation) were simple networks, which have been built upon during the evolution of the more resourceful ones. These primitive solutions used the asymmetry of the network weights to lock onto the ball, where only the difference between the up and bottom sensors of the array implied the direction input to the motor. A very obvious limitation of such networks is that they are unable to actively search, rather remaining at fixed points (i.e., remaining down left until the ball is again in sight) or in a trivial oscillatory behavior. These primitive solutions were gradually substituted by networks that were able to solve the problem more robustly, although aspects of the initial networks have also been inherited by their descendants.

The agents inspected in the following sections have the ability to never lose sight of an object under normal conditions (e.g., unchanged gravity, same size of ball). Many of the networks were not only selected for their high fitness, but also for their observed behavior in different conditions, for example, a smaller ball and higher simulation frequencies (many subtle and interesting properties are hard to define in terms of the fitness function, for example, the search strategies

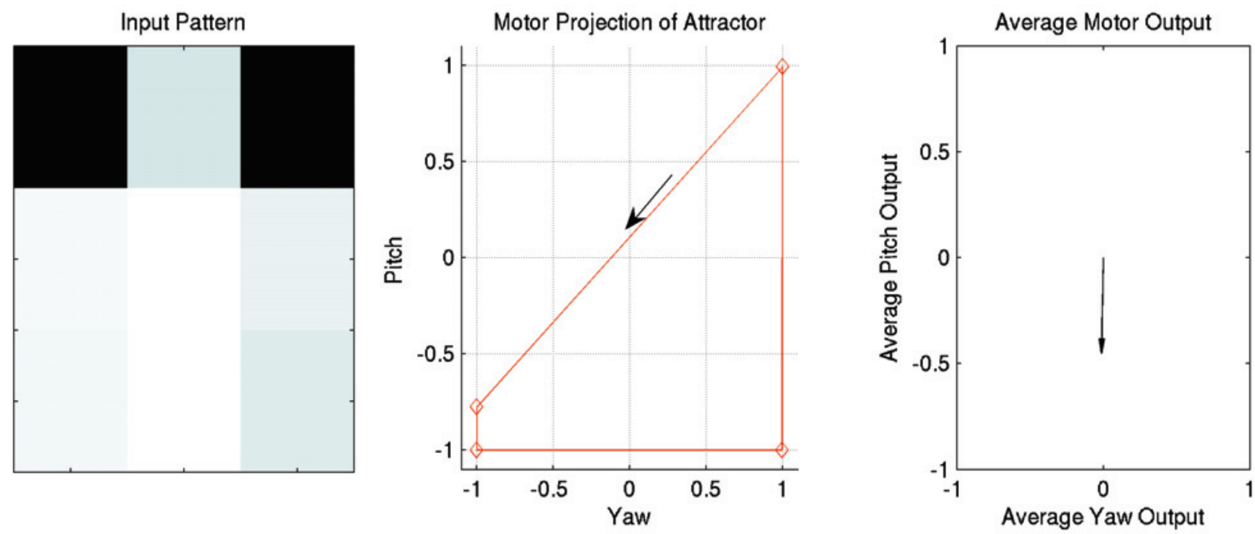

Figure 4 Left to right: stimulus pattern, the period 4 attractor itself, and the averaged output of the associated attractor (the vertices of the polygon are the attractor's states). In the middle box, the arrow indicates the direction of states in the attractor.

when the ball moves out of sight). The networks whose behavior was seemingly less stereotypical also proved to have more diversity of supporting dynamical structures. Nevertheless, the network's size was constrained in evolution to no more than six units in the recurrent layer and a maximal of 120 synapses.

3.3 Analysis of Dynamical Entities Generating Behavior

Most of the analysis is performed in terms of the asymptotic behavior of networks decoupled from the sensory-motor loop of the simulation. Artificial stimulus patterns, which represent interactions of the sensors with the ball, lead to responses that are projected onto motor space. This reduction has often been used in such studies and allows an intuitive understanding of the behavior of the network. It produces a picture of the behavior in the case of constant input, and it depicts the snapshot tendency of the meta-transient. Moreover, to represent the motor actions the agent's body effectively impresses, we average the motor outputs (100 network steps), producing the average amplitude of the motor outputs to verify the action tendency in any given moment.

3.3.1 Velocity Modulation via Motor Projection of Attractors For our analysis, we define the velocity with which the head will move by the average of the output of the motor units for a number of cycles. This is seen in Figure 4, which shows a constant stimulus input, the associated attractor, and an arrow representing the averaged output for both motors.

Figure 5 shows the period 4 attractor in a time series of the motor output projections, in time, for both motors. The output of the motor units is regulated by the profiles of the activation curves for both motors, and therefore by the shape of the motor projections of the attractor. The amplitudes of the motor projections determine the velocity imprinted to the motors.

In Figure 5, we also see that although the activities of the two motors (in any t) are dependent on one and the same attractor, each motor reads different aspects of it. So, although the activity of the network is holistic, the motor efferences are independent readings of this n-D attractor shape, as we have claimed in the introduction.[5]

3.4 Attractor Landscapes

During behavior, the attractors that are projected to motor actions are often not trivial. The network does not simply associate a fixed point attractor in response to a stimulus pattern, such as in a Hopfield network.

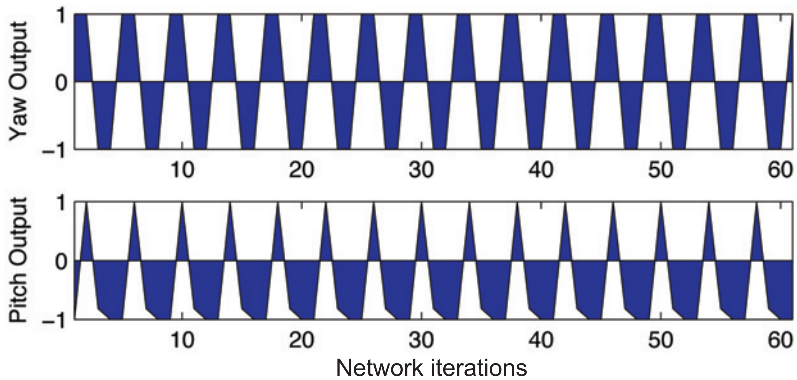

Figure 5 Temporal translation of an orbit on the period 4 attractor, for yaw (top) and pitch motors (bottom). The shape of the oscillations evoked by the attractor defines the velocity arrow in Figure 4.

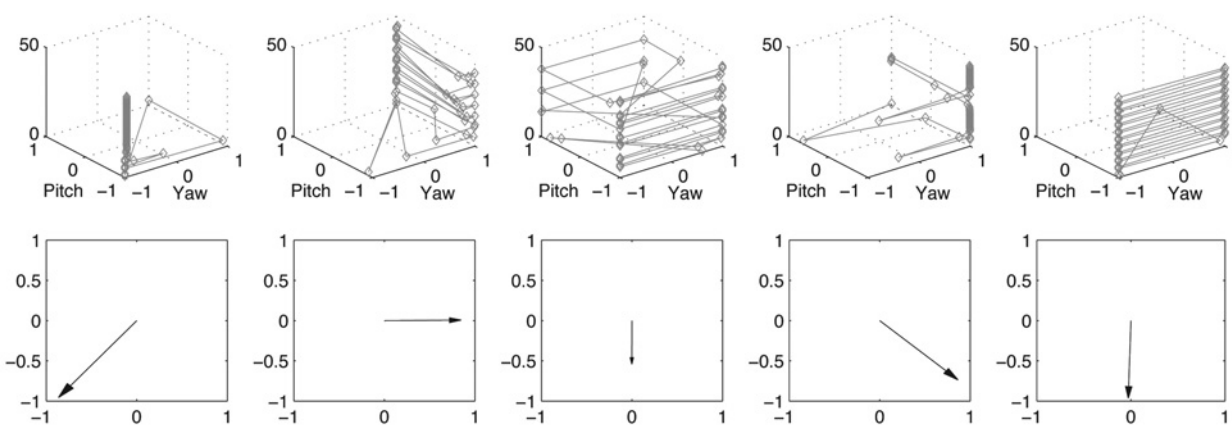

Figure 6 Plots of orbits on output phase space, for a series of different attractors for one and the same input pattern and randomized initial conditions of the hidden layer [relative position of the ball (0,0)]. The *z*-axis tick is the number of network iteration steps. The widely varying shapes of the attractors indicate distinct basins of attraction, and coexisting attractors on the attractor landscape.

In fact, we find a rich portfolio of attractors, as can be seen for example in Figures 6, 8, and 11. We observe that for a large portion of the input patterns (constructed to represent interactions with the ball), the natural asymptotic output of the network is some non-trivial attractor, normally with high periodicity, quasi-periodic orbits or chaos (see Lyapunov exponents in Appendix A). This is also true for coexisting attractors, which exist for one and the same input pattern. This is verified by randomizing the initial conditions of the hidden layer and comparing the resulting asymptotic states for one single stimulus pattern (as in Figure 6).[6] In support of our observations, unlearned randomly connected RNNs with high recurrence have been shown to possess many very informative states (e.g., Berry & Quoy, 2006; Molter, Salihoglu, & Bersini, 2007), as they use many different attractors to translate different input patterns. Here, it is also the case that the networks "freely" associate different types of attractor to stimuli during the evolution, keeping a high capacity for attractor storage.

3.4.1 Implicit Mapping of Environmental Asymmetries It is easy to see that there is no one-to-one mapping of a given input pattern to a velocity of the ball.[7] It follows that to be optimal (gaze locked with the ball), the head has to use different velocities even

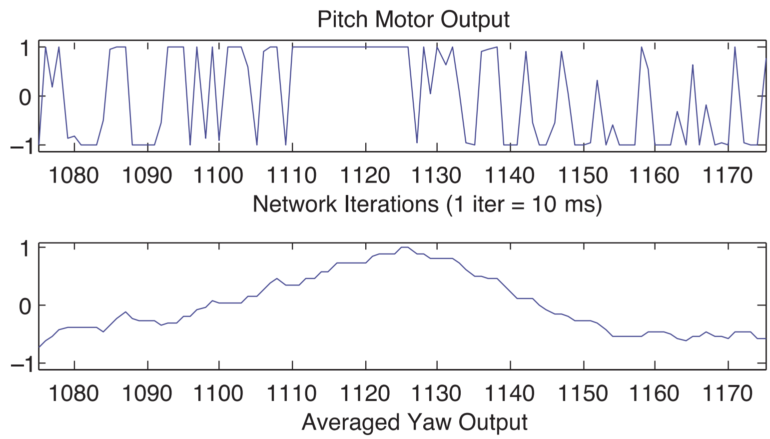

Figure 7 A roughly linear increase and decrease of the pitch velocity of the head. The top row is the actual output of the network. The second row is the convolution with a rectangular causal kernel of 10 steps (0.1 s), representing the average velocity implemented by the tracker. For the average velocity to increase, the meta-transient must be switching to attractors of different shapes across the landscape.

when keeping the input pattern constant (although the acceleration is constant). Therefore, the velocity of the head has to be chosen by considering the recent velocity history in order to respond to gravity. For example, in the case of a ball falling, in order for this to occur the meta-transient has to approach attractors whose motor projection increases as the ball accelerates. Furthermore, as the input patterns are rather similar when the ball is in sight, small paths in parameter space have a very definite meaning.

This is indicated in Figure 7, which shows a plot of the pitch projection of the meta-transient. Here, the activations of the pitch motor unit were recorded during a trial. The rigged profile of the transient is averaged with a causal rectangular convolution window, to represent the effected output velocity. We see that the oscillations on the y-axis lead to a linear increase of the averaged velocity of the tracker's neck. This is consistent with a linear velocity increase imposed by gravity. This means that the network has implicitly imprinted the interaction with gravity into its dynamical substrate. The negative feedback was adjusted to cope with the specific physics of the environment.

3.4.2 Attractor Landscape and Negative Feedback How is the agent able to accompany the ball despite the increasing velocity? Rephrasing this question in terms of the tendencies under which the meta-transient is subject might give us some insight into the mechanics of control. By plotting the squashed projection of the output space as a function of the possible interactions of the input array and the ball, we gain an outlook of the whole of the agent's action set.

For each output unit, we plot the mean amplitude of the respective motor projection for all the states in input space, as in Figure 8. The figure is constructed as follows. The coordinates of each pixel denote the relative position of the ball in head-centered coordinates. The color of this pixel represents the average amplitude of the respective output unit given one dynamical system, parametrized by the correspondent interaction with the ball, in head-centered coordinates. The center of the diagram (0, 0) denotes when the center of the ball image coincides with the center of the retina; that is, the relative position of the ball to the retina, both in horizontal and vertical coordinates, is zero (to see how the stimulus looks at the center of the coordinate system, check the last state in the input sequence in Figure 3, which is about the center). To compute the action to every pixel, we proceed systematically by calculating the actions associated with the pixels according to vertical scan lines. For every new input pattern, we then calculate 300 steps in that orbit, drop 200, and average the last hundred.[8]

Here, we describe how Figure 8 explains control. For each stimulus (i.e., for each given initial condition resulting from interaction between head and ball), there is a single pixel that represents it, for each associated output unit. Assume that for a small number of

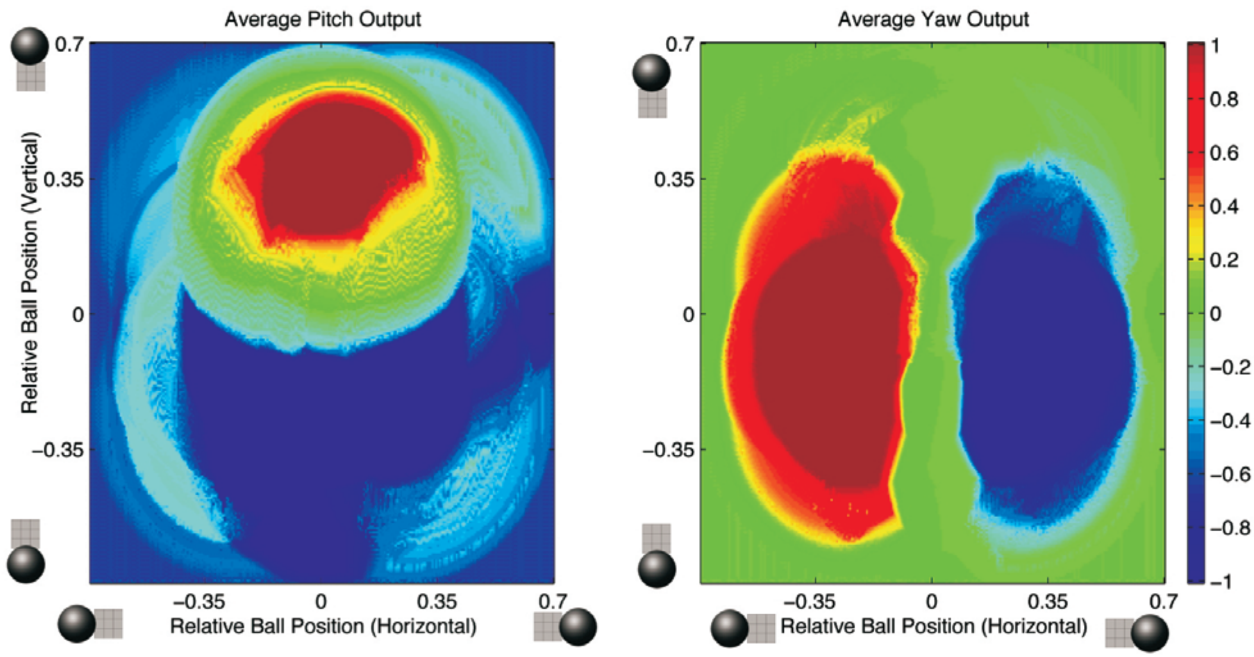

Figure 8 Depiction of the attractor landscape for the totality of the possible input space, on the left the pitch motor unit, and on the right the yaw motor unit. Every pixel is a head-centered coordinate of the ball (the icons in the corners represent the interaction between the input array and the ball). The behavior of the agent can be seen as the meta-transient relating paths in parameter space (see text for explanation).

network iterations, the input remains similar. The color of that pixel then represents roughly the action of the agent, albeit not exactly, as it is dependent on a temporal average, and the meta-transient might become entrained in different states of the attractor. The simplest example is as follows. When the ball is in the precise (0, 0) coordinate, the action of both the yaw and pitch motors is roughly zero. Assume the ball is falling. If the agent does nothing, the sensory input will change, and the ball will be lower in relation to the center of the retina [e.g., coordinate (0, –0.3)]. This new situation will remove the agent from repose and evoke a response of the pitch motor close to –1, while the yaw remains at zero. However, because –1 is the largest velocity, the head will advance in comparison with the ball, for example changing the coordinate to (0, 0.3), which will evoke a small upward velocity. This process continues as connected paths of parameter space (stimuli) are associated with attractors (motor actions). It should be clear that the actions of the agent result from a negative feedback loop, where the environment provides the control signal. Gravity is encoded in the shape of the loose input-action mapping of the attractor landscape, and expressed in the sequence of states of the meta-transient.

It can be said that Figure 8 depicts the agent's "action space," by showing what the velocity imprinted by the motors would be, had the input been one given input pattern, when the network is given enough time. However, recalling that the ball dynamics does not permit the network to settle on its final attractor lazily, the picture should be read as a collection of input-tendency pairs. Without enough time to settle on the attractor, the behavior is the meta-transient that overlays the attractor landscape, as a function of the connected path in sensor/parameter space.

It is also interesting to observe that the functional characteristics of the individual projections on the output dimensions (pitch and yaw) are strongly dissimilar, although both are functions of the high-dimensional space of the evolved network. Conceivably, this results from the impact of the asymmetries of the environment and of the ball's behavior: in the vertical direction, a bouncing driven by gravity, while in the horizontal direction, a more constant velocity (contingent on the wall bounces).

These landscapes varied considerably in the details among the evolved networks, although the general characteristics were present in all the successful networks. That is, despite the randomization of the different evolutionary runs that produced the networks, the evolutionary path generated attractor landscapes that would provide for a similar function. Moreover, the general trend in evolution was one of complexification of the landscapes: increased complexity of the underlying attractor landscapes usually implied solutions that were both more robust and resourceful (an aperiodic active search, for instance).

3.5 Attractor Shapes and Action

The averaged motor projections of the attractors also indicate that there are two features of the attractors responsible for the action, where neither is preponderant over the other (i.e., the attractor's periodicity and shape). Different periods might lead to similar average speed, and conversely equal periods might lead to different outputs. This is clearly illustrated with the following series of chaotic and quasi-periodic attractors in Figure 4, calculated for very similar input patterns.[9] For the Lyapunov exponents of the respective stimuli, please see Appendix A. Figure 4 shows the motor space projection of the simulated asymptotic orbit on the attractor, for 150 steps, with similar averaged velocities. In the simulation, however, there would probably not be such equality. As the actions during the trials occur over a very short time, the average speed might change depending on two factors: where the transient starts and for how long it approaches the attractor. This is particularly true for chaotic attractors, in which transients might become entrained in different positions. From Figure 9 we see that the attractor has a definite shape and structure. However, because it is aperiodic, the average over motor output will obviously vary depending on the time window taken for computation. Nevertheless, it will only vary within the

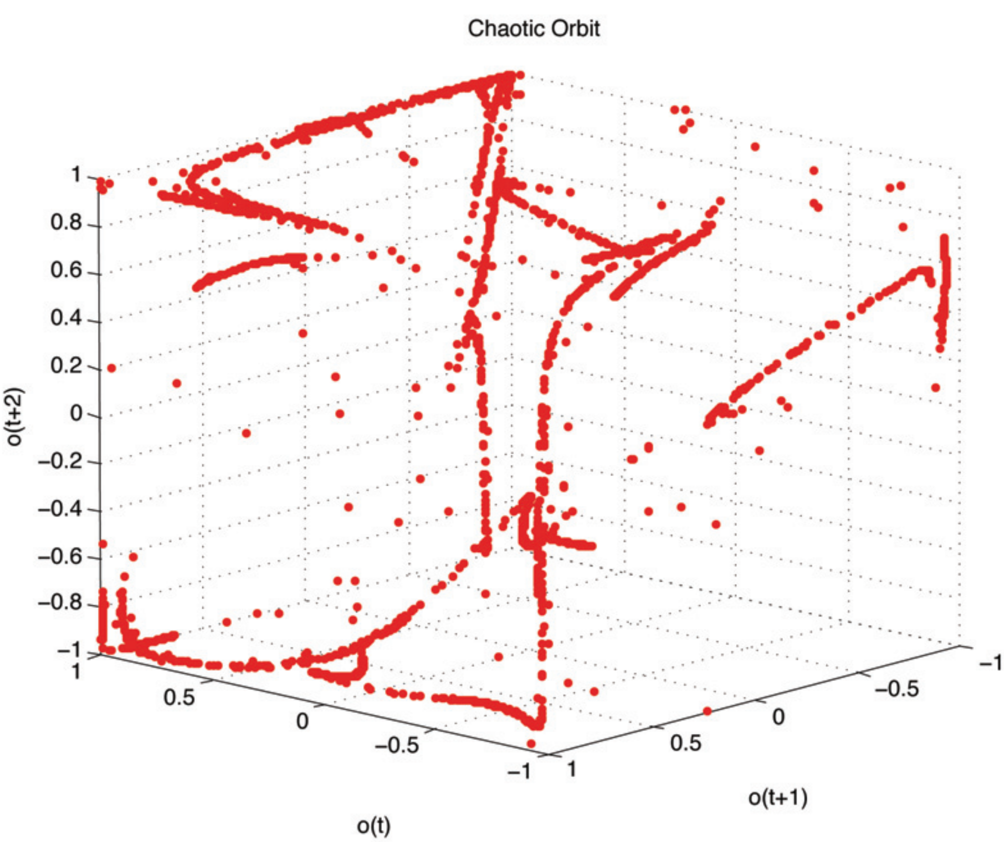

Figure 9 The first-return map of a chaotic attractor for the yaw motor output indicates that despite the presence of structure, the putative action of the agent will change depending on where the meta-transient engages the attractor.

bounds given by the structure of the attractor itself (i.e., depending on where the meta-transient engages the attractor, and how long it remains under its influence).

3.5.1 Attractor Morph The attractors might also change smoothly with paths in parameter space contained within bifurcation boundaries. A meta-transient guided through a morphing attractor might wind along the attractor set, changing smoothly as the attractor landscape is explored through parameter shifts. Figure 10 shows a morphing attractor given a particular path in parameter space. Basically, the plot is the exploded version of one vertical scan line of Figure 8, where the output is not squashed in the average. The z-axis represents the parameter path, corresponding to the sensory input from the ball on the bottom of the sensor to the ball on top. In other words, the z-axis in the figure represents the relative position of the ball in the vertical direction, keeping the horizontal constant. The x- and y-axes are the motor phase space. Each slice in z determines a dynamical system given the corresponding parametrization (stimulus). For each slice, we plot 45 steps, after discarding the first five, of the projection of the orbit. The plot is in fact a bifurcation diagram, where we plot the two output dimensions as a function of the parameter, which is the relative position of the ball.

We interpret Figure 10 as follows. We start, for example, at the bottom of the z-axis with the relative ball position on the vertical axis = –0.7 m (the relative position on the horizontal axis is constant at –0.23 m). In this situation, with the ball having a slight offset to the left, the network starts with a saturated response of –1 in yaw and –1 in pitch, and very low amplitude oscillations. As the parameter in z is increased, meaning the ball moves up relative to the input array, the amplitude of

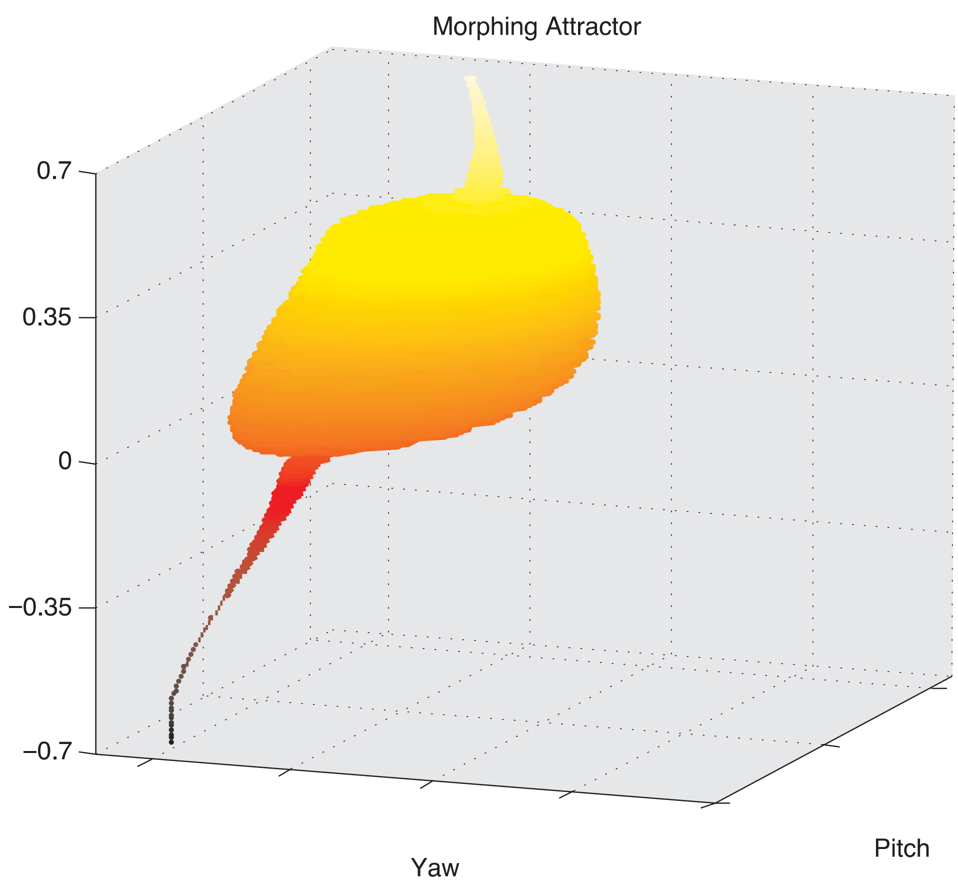

Figure 10 Morphing attractor for a path in parameter space corresponding to a ball's trajectory from bottom to top (z-axis). For each parameter in z (the relative position of the ball on the vertical axis), the orbit is computed in phase space and plotted as a z slice.

the oscillations on the phase space increases smoothly. The density and form indicate that the morphing attractor might be composed of quasi-periodic attractors. As we continue to move up the z-axis (i.e., ball up), the attractor smoothly saturates again on the maximum pitch. Because these transitions are between bifurcation boundaries, the shape of the attractor set is smooth. It is also worth mentioning that this particular solution was not the case for all the networks, despite their apparent similarity in action sets.

3.5.2 Coexisting Attractors The presence of coexisting attractors is observed by presenting the same stimulus but with randomized initial states of the hidden layer. In theory, two identical parametrizations possess the same attractor structure, because they resolve the same dynamical system. Theoretically, the fact that two initial conditions lead to different attractors means that there are at least two attractors, each belonging to a distinct basin of attraction, possibly across an unstable manifold. When we observe different orbits with the input pattern clamped constant (a single dynamical system subject to one parametrization $\mathfrak{D}_{\rho_{\text{pattern}}}$), we can conclude that there are coexisting attractors for the dynamical system determined by that parametrization. Coexisting attractors might play similar or different functional roles, as discussed later.

In Figure 6 we see that the network might produce distinct outputs when given one and the same stimulus pattern, plus a randomized initial condition of the hidden layer. This indicates the coexistence of attractors with different functional roles, because they produce different average motor outputs as the arrows point in different directions.

We compare Figure 11, where we plot the actual attractors elicited during the trials (the activities of the hidden layer were logged, allowing off-line recomputing of the attractors visited during the trials). The stimulus was the same for all the figures, but the initial condition of the hidden layer was taken from the trials. The z-axis is the iteration step of the network, so we can see the time series of the motor projection. Note that, although the shape of the attractor noticeably varies, the average outputs of the attractor (given by the arrow plot beneath) are in this case very similar. This indicates distinct coexisting attractors, which nevertheless play similar functional roles.

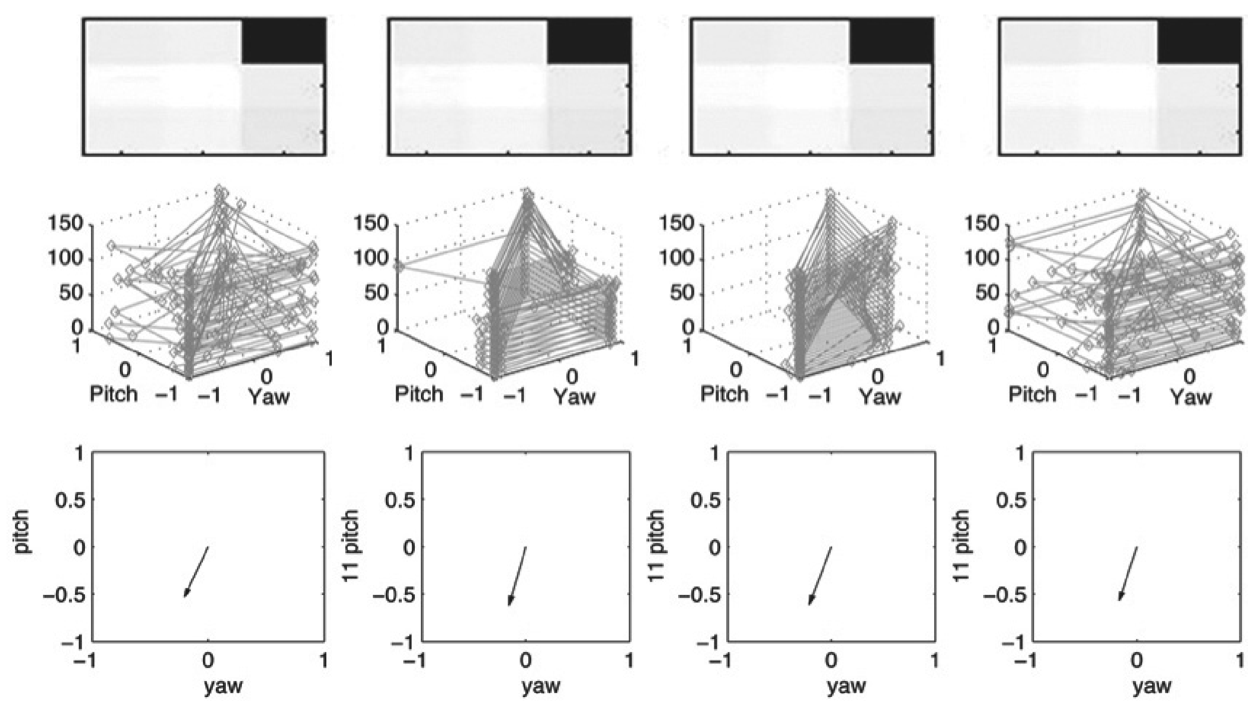

Figure 11 For each of the stimulus patterns, the motor output orbit is plotted in phase space (the z-axis is the iteration cycle). Although five different attractors appear, they all lead to very similar actions (compare the velocity vectors of the third row).

Because the coexisting attractors are accessed in accordance with the internal state of the hidden layer, this invokes high context dependency. As the system is sensitive to history, this implies a sort of transient memory. Transient memory is an issue of current cognitive dynamics research (as in, for example, Beer, 2008). The fact that there are coexisting attractors implies that there are multistable states, and so different actions can be elicited by one and the same stimulus as a function of history.

4 Discussion

4.1 Transmission of Attractor Projections: Modularity

We have shown that the motor projection of the meta-transient is the functional underpinning of the competent behavior that the agent deploys. However, as discussed in the introduction, and seen in the phenomenon of morphing, the meta-transient can be said to draw shapes in phase space, and also in the motor projections thereof. Therefore, we may say that the shape of the attractor, at least in the case shown, has functional significance, inasmuch as it is what determines the translation from shape to actions in order to command the motors.

If it is correct to employ attractor shapes as one explanatory feature of the theory of cognition, other promising research fronts appear, such as, for example, how attractors might communicate their meaning (i.e., their identifying shapes) through transmission bottlenecks (such as axonal pathways) to neighboring modules. Here, modules (Pasemann, 1995) are taken as the brain's computational units according to some criterion, which take afferent inputs and produce efferent outputs.[10,11] In this article, what modules project to others is not the abstract high-dimensional space activity, but much simplified projections, with retrievable coding schemes.

Identifying how transmissions carry the important features of the attractors to and from other modules, such as for example hyper-columns or back projections to the sensorial preprocessing, may produce insights into how to evolve functionality modularly. In other words, an approach based on the attractor identities projected to smaller dimensions might give insights into how to define neural interfaces of dynamical exchange. However, in holistic systems, the embodied brain attractors are never independent of each other, because of massive efferent and afferent connectivity between brain areas and modules, but all the same, the messages they carry across can only be a sample of the internal activity of these modules. A requirement for pursuing this idea of modularity further is whether it is at all feasible to disentangle the attractors or whether it might be a lost cause (Watson & Pollack, 2005). Nevertheless, many do think that modularity in the brain is a pervasive characteristic (Fodor, 1983), interfaces remaining to be defined (Haugeland, 1995).

4.2 Learning as Deforming Attractor Landscapes

Neurodynamics defines learning as changes of the network during behavior (i.e., variation of weights according to synaptic plasticity rules). The networks here were static; plasticity was not involved. This simplification, which permits simpler off-line analysis, has the disadvantage that all of the landscapes have a rigid structure. In spite of this, active tracking as defined here is a simple problem, and one that can be solved merely with evolution. The selection process can be seen as operating on attractor landscapes, selecting those with the most useful action set.

In living beings, ecological problems are way beyond the isolated function of tracking. One of the essentials is learning. However, thinking in terms of the dynamical substrates of the embodied RNN provides a new facet to conceptualizing learning. We might say that evolution begets initial parameter sets (instinctive responses), while learning modulates them to extend behavioral breadth (Sterelny, 2005). Stretching out the analogy to the development of the central nervous system of higher organisms, it could be said that the initial structure has to carry those sets of conditions that allow for the lodging of the "learnable" attractors, those that the organism is able to learn. This provides a neurodynamics interpretation of the Baldwin effect in terms of the "learning capacity" of the underlying networks. Selection will benefit those individuals endowed with initial structures that are more malleable and permeable to useful attractors landscapes.

4.3 Related Work

There has been much work carried out over the last 15 years on the topic of the dynamical systems analysis of evolved agents, pioneered in articles by, for example, Husbands, Harvey, and Cliff (1993), Beer (1995), Tani (1998), and others. Our contribution rests on the shoulder of these seminal works, and presents a novel toy problem on pan–tilt tracking, with high environment–sensory-motor dynamics and an analysis of behavior based on evolved attractor landscapes (rather than learned, as in for example Ijspeert, Nakanishi, & Schaal, 2003), which gain a little more substance than in the usual metaphorical usage (e.g., Banerjee, 2001).

With reference to the concepts of the meta-transient and attractor landscapes, the concept of chaotic itinerancy (CI) advanced by Ikeda, Kaneko, Tsuda, and others is worth mentioning for two reasons. First, this has made great impact in the field of cognition and brain dynamics (Tsuda, 1991, 2001), not only very broadly (Kaneko & Tsuda, 2003), but also formally (Kaneko, 1990), conceptually (Tsuda, 1991), and metaphorically. An excellent review of the applications of the concept is presented in Kaneko and Tsuda (2003). Second, important differences appear when the approaches are contrasted. Mercilessly squeezing the concept into a couple of sentences, chaotic itinerancy is the idea that a chaotic attractor spanning the whole of a high-dimensional phase space of a dynamical system (also neural networks) at times collapses in limit cycles of lower dimensions, called attractor ruins. The paths connecting the lower-dimensional attractor ruins are itinerant (chaotic) because of the crossing of unstable manifolds, thus giving the name chaotic itinerancy.

Although the concepts of chaotic itinerancy and meta-transient can be made analogous, there are two main distinctions to be drawn. First, the systems we describe are parametrized by the input, and therefore we deal with a collection of dynamical systems, instead of a single one with many dimensions. So, in the case of itinerancy we talk about one orbit that explores the complexity of the basins of attraction of one dynamical system of high dimension. However, here we talk about the complexity of the attractor landscape, which is the concoction of all basins of attractions of all dynamical systems in an RNN accessible by parametrizations. Second, in chaotic itinerancy the basin crossing of an orbit is usually due to noise, and unstable manifolds, while here the meta-transient crosses dynamical systems because of parametrizations of the system, which are a function of the structural coupling between the agent and the environment. So, the crucial aspect of this difference is parametrization by structural coupling. Ikegami and Tani (2001) touched on the core of this issue in the title of their article: "Chaotic itinerancy needs embodied cognition to explain memory dynamics." Finally, because our functional states are efferent projections, the actual dimensionality of the attractor in higher dimensions is not crucial for behavior; there is a many-to-one mapping of different attractors to actions. It is assumed to be indifferent whether the period of the attractor is low or high, as long as the motor projection finishes the job.

5 Conclusion

In their renowned article, von Holst and Mittelstaedt (1950) wrote about the reflex arc: "…jede Einzelbewegung reflektorisch ihre Gegenbewegung …in Gang setzt."[12]

Introducing some controversy by interpreting the model presented here as an analogy, it is possible to frame the phrase above in terms of behavior on the sensory-motor loop, as a sequence of non-trivial reflex arcs, with reafferences from the state of the body, the control structures, and the environment. This is because for every state of a path in parameter space, and every initial condition of the hidden layer, there corresponds an attractor, whose mapping to action causes the next state of the path in parameter space, closing the loop in a sequence of linked actions. However, this does not necessarily mean that reflexes are stereotypical.[13] The complexity of attractor landscapes and possible coexisting attractors may lead to a wide behavioral breadth (Sterelny, 2005), so that complex reflex arcs might generate highly adaptive behavior, through context dependency, redundancy, and novelty in response. Here, these three are exemplified through coexisting attractors, different attractors with similar meanings, and chaotic attractors entrained in different moments. Neurodynamics has provided the analysis methods for the evolved embodied RNN, where the concepts of attractor landscapes and meta-transients show the RNN to be rich in behavior capacity.

In the context of solving problems of physical nature, we have studied a simulation that has shown how an evolved attractor landscape provides for embodied active tracking. This comes from the hypothesis that motor control, because it is dynamics (of body) embedded in dynamics (of environment), can be properly described and analyzed using the language of neurodynamics.

5.1 Future Work

We are currently running a variety of experiments that extend the experiment described here and probe the dynamics generated by adding proprioceptive inputs to the network. Specifically, one experiment removes the restriction of the ball only moving in one plane, by allowing the ball to bounce in a room with three dimensions. Preliminary results indicate that despite the additional source of ambiguity of ball proximity, tracking is successful. Additionally, we are studying the dynamical mechanisms of distance estimation by vergence, by introducing a second pan–tilt and postulating as fitness the correlations between the input arrays. For both the above-mentioned experiments, the analysis will derive and extend the analysis in this article, by outlining the importance of the embodiment variables (internal feedback through proprioception), and demonstrating the space of possible dynamics of interactive behavior buttressed on the concepts of attractor landscapes and meta-transients.

Appendix A: Lyapunov Landscape

In many parts of this article, we have claimed that there were chaotic attractors in regions of ambiguity and that even chaotic attractors with positive Lyapunovs may be functional states. In support of this, we plot the landscape of Lyapunov exponents for the different possible parametrizations of the network (stimulus space) using the same axis dimensions as in the attractor landscape (Figure 8). For each pixel on the plot, the state of the hidden layer is randomized. In this way we also sample the coexisting attractors involved.

Two aspects should be observed. First, as the ball comes into the center view, there is a tendency for the attractors to have positive Lyapunov exponents. We attribute this to regions of larger ambiguity (is the ball coming into or leaving sight?). Second, the scattering of points of different Lyapunov values spread over the landscape indicates that for that particular dynamical system we calculate the Lyapunov of different attractors; in these regions there may be the coexistence of chaotic attractors and attractors of different periodicities.

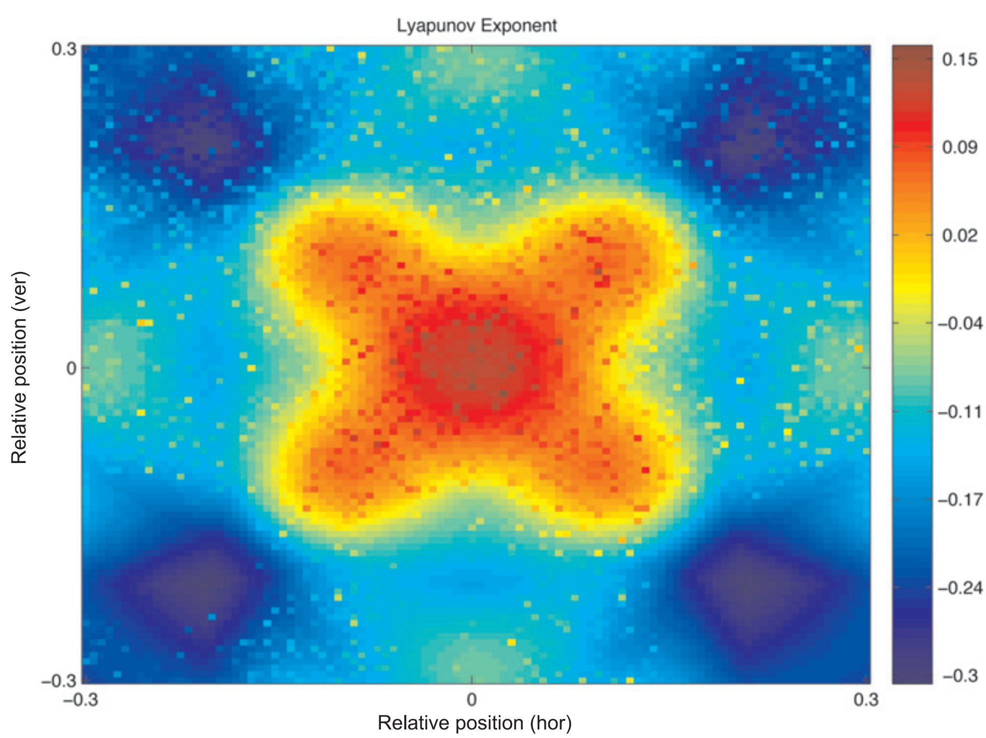

Figure A1 Each pixel represents the Lyapunov exponent of one of the attractors (the initial condition of the hidden layer is randomized) for each of the stimulus patterns. The *x*- and *y*-axes are the position of the ball in head-centered coordinates.

Acknowledgments

This work was partly supported by the DFG grant CH 749-1. The authors would like to thank the three reviewers, whose clear and precise comments enhanced the transparency and readability of the article.

Notes

1 We use "parametrizations" here in the manner that has become usual in the literature of dynamical systems, in which parameters are variables that change slowly with respect to the dynamics of the network.

2 This is meant in the sense that not all four-sided figures are squares and not all squares are of the same size. In other words, the number of edges does not exhaust the relevant aspects of the figure. The statement about attractor shape identity suggests that attractors can be seen as figures in phase space, which is, after all, a measurable space. In this sense, two period 5 attractors occupying different portions of the phase space are not the same if they do not look the same.
An orbit draws paths that are constrained, or guided, by the influence of the attractor whose basin it occupies. In this sense, even chaotic attractors, with their ineluctable difficulties, may be described by their shape. By considering the fact that chaotic attractors cover regions of phase space with varying density, this helps to illustrate that as an orbit approaches the attractor, it outlines its shape. This is like discovering a solid when blindfolded: although we never touch the whole solid at a time, as we continue to touch it, its shape is incrementally revealed to us.
We believe that this metaphorical notion becomes useful in our discussion of attractor morphing later, in which we catch glimpses of the shape of the N-dimensional attractor by seeing how its lower-dimensional projections change as we parametrize the network. In fact, attractor shapes become most meaningful in a discussion of the functional roles of the activities of a network, as here, where the amplitude of an oscillation is translated to the velocity of the motor, subject to the inertial masses involved, or in other words, the projections of the attractor's shape do help to shape the agent's actions.
Although we are not aware of any particular reference in the literature concerning the functional significance of particular attractor shapes, the distinctive geometrical features of the attractors commonly studied in the literature (e.g., limit cycles of planetary orbits or the mandelbrot set in the study of biological morphology) often define (or strongly correlate with) the behavior of the system they describe. Therefore, we have the idea that an attractor's shape is its identity.

3 It is helpful to be aware that "qualitatively distinct dynamic behavior" might lead to qualitatively similar agent behavior, which might be a source of ambiguity. Distinct dynamic behavior refers only to domains of different attractors, while qualitatively distinct agent behavior is a much more subjective criterion, requiring the judgment of an observer.

4 The weight matrix and bias vectors of the network in Figure 1:

$$W = \begin{pmatrix} -18 & 8 \\ -8 & 0 \end{pmatrix} \quad \text{and} \quad \Theta = \begin{pmatrix} -0.45 \\ \theta_2 \end{pmatrix}.$$

5 The idea that readout units sample different aspects of the same attractor is analogous to the approach of liquid state machines or echo state networks (Jaegger, Maas, & Markram, 2007). The difference is that their attractors are generated randomly to satisfy certain requirements for complex dynamics, while our attractors are incrementally evolved. This difference results from the role of the attractors, in which we think that there are attractors that are more apt to solve some types of problem, and that artificial evolution is a good method to beget them.

6 This does not mean, however, that all the agents reach every coexisting attractor during behavior, as the possible states are also bounded by the possible history of interactions.

7 Consider that when the input pattern is locked with the ball, the input pattern is constant, but the pitch motor has to accelerate down.

8 The internal states of the hidden layer are inherited across parametrizations.

9 These input patterns were obtained with the same network operating at 500 Hz update frequency, so the differences between steps would be tiny. Note that this test would not have been possible if the network had not been robust to different update frequencies.

10 Arguably, it might not be the case that the computational units can be easily outlined and modularized. However, there are extensive data showing the presence of possible interfaces in brain structures, given their observable organization of architectonic implementations. Spatial spread incurs in time specific operations, resulting from axonal delays and distinctive connectivities. At one extreme, Fodor (1983) argues for absolute modularity, while Haugeland (1995) states that there is no possible unequivocal unpluggability criterion; these extreme positions are held by many others (see for example, Clark & Chalmers, 1998; Grush, 2004). Still others (such as Freeman, 1995) argue that although the resulting activations are necessarily holistic and hermeneutic, different mod-

ules contribute in specific ways to the global patterns of global significance. Whatever the case, if we aim for agreement with the neuroscientific understanding of brain function, as modelers we should propose mechanisms of information exchange that will allow a certain degree of localization. Here we take the semirealist approach, where although the processing is happening on the higher-dimensional space of the computing unit, the projection of the attractor carries orderly correlates, and thus products that can be safely used by a number of modules of different roles. Besides, to take an attitude of mystery is to forsake many glimpses of understanding that do seem to point out that the brain might rely on an abstract and general neural mechanism, which is locally implemented.

11 Information and computation here are used with small "i" and small "c."

12 …that each movement reflexively starts its counter-movement and so, in a sequence.

13 Merleau-Ponty (1965) made a case for the inadequacy of a physiological theory of behavior based on reflexes, as a quantitative change of the stimulus induces a qualitative change in behavior. However, as we have seen in our case, even a small quantitative difference does precisely this, it invokes a reaction that is incommensurable with the stimulus, but adequate to solve the postulated problem.

References

Banerjee, A. (2001). The roles played by external input and synaptic modulations in the dynamics of neuronal systems. *Behavioral and Brain Sciences*, *24*(5), 811–812.

Beer, R. (1995). A dynamical systems perspective on agent–environment interaction. *Artificial Intelligence*, *72*(1), 173–215.

Beer, R. D. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, *18*, 3009–3051.

Beer, R. (2008). Beyond control: The dynamics of brain–body–environment interaction in motor systems. In D. Sternad (Ed.), *Progress in motor control: A multidisciplinary perspective*. Berlin: Springer.

Beer, R., & Gallagher, J. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, *1*(1), 91–122.

Berry, H., & Quoy, M. (2006). Structure and dynamics of random recurrent neural networks. *Adaptive Behavior*, *14*(2), 129–137.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.

Dieckmann, U. (1995). Coevolution as an autonomous learning strategy for neuromodules. In H. J. Herrmann, D. Wolf, & E. Pöppel (Eds.), *Supercomputing in brain research: From tomography to neural networks* (pp. 427–432). Singapore: World Scientific.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Freeman, W. (1995). *Society of brains: A study in the neuroscience of love and hate*. Philadelphia, PA: Lawrence Erlbaum.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery and perception. *Behavioral and Brain Sciences*, *27*, 377–442.

Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, *11*(1–2), 79–98.

Haugeland, J. (1995). Mind embodied and embedded. *Acta Philosophica Fennica*, *58*, 233–267.

Huelse, M., Wischmann, S., & Pasemann, F. (2004). Structure and function of evolved neuro-controllers for autonomous robots. *Connection Science*, *16*(4), 249–266.

Husbands, P., Harvey, I., & Cliff, D. (1993). Analysing recurrent dynamical networks evolved for robot control. In *Proceedings of the 3rd IEE International Conference on Artificial Neural Networks (ANN93)* (pp. 158–162). London: IEE Press.

Ikegami, T., & Tani, J. (2001). Chaotic itinerancy needs embodied cognition to explain memory dynamics. *Behavioral and Brain Sciences*, *24*(5), 818–819.

Ijspeert, A. J., Nakanishi, J., & Schaal, S. (2003). Learning attractor landscapes for learning motor primitives. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems 15*. Cambridge, MA: MIT Press.

Jaegger, H., Maas, W., & Markram, H. (2007). 2007 special issue: Echo state networks and liquid state machines. *Neural Networks*, *20*(3), 290–297.

Kaneko, K. (1990). Clustering, coding, switching, hierarchical ordering, and control in network of chaotic elements. *Physica D*, *41*, 137–172.

Kaneko, K., & Tsuda, I. (2003). Chaotic itinerancy. *Chaos*, *13*(3), 926–936.

Merleau-Ponty, M. (1965). *The structure of behavior*. London: Methuen.

Molter, C., Salihoglu, U., & Bersini, H. (2007). The road to chaos by time-asymmetric Hebbian learning in recurrent neural networks. *Neural Computation*, *19*, 80–110.

Negrello, M., & Pasemann, F. (2008). Adaptive neurodynamics. In Y. Shan & A. Yang (Eds.), *Applications of complex adaptive systems*. Hershey, PA: IGI Publishing.

Pasemann, F. (1995). Neuromodules: A dynamical systems approach to brain modelling. In H. J. Herrmann, D. Wolf, & E. Pöppel (Eds.), *Supercomputing in brain research: From tomography to neural networks* (pp. 331–347). Singapore: World Scientific.

Pasemann, F. (2002). Complex dynamics and the structure of small neural networks. *Network: Computation in Neural Systems*, *13*(2), 195–216.

Pasemann, F., Steinmetz, U., Hülse, M., & Lara, B. (2001). Robot control and the evolution of modular neurodynamics. *Theory in Biosciences*, *120*, 311–326.

Philipona, D., O'Regan, J. K, Nadal, J.-P., & Coenen, O. J.-M. D. (2003). Perception of the structure of the physical world using unknown multimodal sensors and effectors. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems 16*. Cambridge, MA: MIT Press.

Simon, H. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.

Sterelny, K. (2005). *Thought in a hostile world*. Cambridge, MA: MIT Press.

Tani, J. (1998). An interpretation of the "self" from the dynamical systems perspective: A constructivist approach. *Journal of Consciousness Studies*, *5*(5–6), 516–542.

Tsuda, I. (1991). Chaotic itinerancy as a dynamical basis of hermeneutics in brain and mind. In *Microcomputers and attention*. Manchester, UK: Manchester University Press.

Tsuda, I. (2001). Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behavioral and Brain Sciences*, *24*, 793–847.

Varela, F. (1979). *Principles of biological autonomy*. Amsterdam: North-Holland.

von Holst, V. E., & Mittelstaedt, H. (1950). Das reafferenzprinzip (wechselwirkung zwischen zentralnervensystem und peripherie). *Naturwissenschaften*, *37*, 464–476.

Watson, R. A., & Pollack, J. B. (2005). Modular interdependency in complex dynamical systems. *Artificial Life*, *11*(4), 445–457.

About the Authors

Mario Negrello is a mechanical engineer holding an M.Sc. in cognitive science from the University of Osnabrück, Germany (2006). With his Ph.D. underway, he holds the position of an associate researcher at the Frauhofer Institute, Germany, and is currently a visiting researcher at Hod Lipson's Computational Synthesis Lab at Cornell University, made possible through and award from the International Neural Network Society (INNS). As an advocate of the prominence of interaction of the nervous system and body morphology to shape thought and behavior, he continues to pursue neurodynamics as a fertile model for the study of cognition.

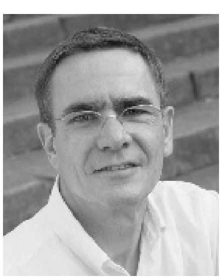

Frank Pasemann studied physics and mathematics at the Universities of Marburg and Würzburg, and at the International Center for Theoretical Physics (ICTP), Trieste: diploma (1971), Dr.rer.nat. (1977) and Habilitation (1985) in theoretical physics. Since 1985, he has been Univ.-Professor, since 1992 Apl.Professor for theoretical physics at Technical University Clausthal, and since 2004 Hon.Professor, Institute of Cognitive Science, University Osnabrück. He has been the head of the research groups at the Research Centre Jülich (1993–1996), the Max-Planck-Institute for Mathematics in the Sciences, Leipzig (1997–1999), the University Jena (2000–2001), and at the Fraunhofer Institute IAIS, Sankt Augustin (since 2002). He is also with the Institute of Cognitive Science, University of Osnabrück, Germany. *E-mail*: frank.pasemann@iais.fraunhofer.de