

Chapter 10

Action Switching in Brain-Body-Environment Systems

Eran Agmon

10.1 Introduction

In recent years, the cognitive sciences have been converging upon an integrated perspective, a perspective that reframes behavior and cognition as a special type of self-organization that arises through the nonlinear, distributed interactions between brain, body and environment (abbreviated BBE). The BBE perspective has been separately developed by multiple lines of research such as the extended mind (Clark and Chalmers 1998), distributed cognition (Hutchins 2000), embodied cognition (Clark 1998), enactive cognition (Noë 2005; Thompson 2007; Varela et al. 1992), situated cognition (Clancey 1997; Hutchins 1995), and the dynamical approaches to cognition (Beer 1995b; Thelen and Smith 1996; Kelso 1995; Port and van Gelder 1995). These different theories all emphasize different elements of the BBE; either the body, or the environment, or the temporal element. But their different theories are friendly to each other and can be brought together into a broader, integrated perspective. By bringing focus to all of the relevant components and their interactions, cognitive systems are transformed into seemingly self-organizing systems, in which behavior and cognition become a dynamical process that unfolds through distributed interactions (Kelso 1995; Maturana and Varela 1980; Thompson 2007).

We need now proceed with caution; the term self-organization has been a widely used term in scientific fields from physics to human social networks, and has acted as a unifying theme in systems sciences such as cybernetics and complex systems. But when we consider the different definitions surrounding self-organization, there is an abundance of philosophical stances and formal methodologies (Polani 2008). In this chapter we will focus on a single perspective of self-organization, which will be called the “absolute system” perspective, after Ross Ashby’s framework for describing what he considers to be adaptive behavior and self-organization (Ashby 1962, 1952). An absolute system takes what some describe as the omniscient perspective (Dupuy 2009), which takes all of a system’s relevant variables and puts

Eran Agmon

Indiana University, 107 S. Indiana Avenue, Bloomington, Indiana 47405-7000, United States
e-mail: agmon.eran@gmail.com

them in a model that fully describes the system's dynamics. For Ashby, a system has a finite set of internal states and a transformation rule that maps a state onto itself as it unfolds in time. An absolute system is an autonomous system, in which all relevant variables are accounted for. With such systems an initial state has a regular trajectory that follows, and upon repeated re-initializations to the same initial state there are no divergences in following behaviors. If there are divergences, then some relevant variables must not have been accounted for and the system is not absolute. A characterization of systems in this sense is ideal for science, because it allows for perfect predictability of the system's behavior. It can be argued that striving for an absolute system description is not practical when dealing with real-world systems for which there are essentially infinite relevant variables. But in this chapter we follow through with the assumption to see what insight can be gained.

Sect 2 describes this perspective in greater detail, and follows Ashby's argument, which he believed demystified the notion of self-organization by attributing the apparent self-organization to an opportune matching of system and environment. We move along this intellectual thread leading from Ashby's definition of self-organization to the modern theories of the BBE framework and show that a very similar approach has been converged upon, perhaps unknowingly, by an integrated BBE perspective. We will see that Ashby's insights are pervasive in this framework, and have been developed into to a rigorous research methodology. The desire to understand the relevant causal variables that come together to generate our cognitive behavior has led scientists to "extend" the mind (Clark and Chalmers 1998), and to attribute mental processes to the complex interactions that take place between many distributed components in the brain, body, and environment. The field is developing a terminology much like Ashby's, which is rooted in dynamical systems theory, and focuses on the temporal elements of cognitive behaviors that arise through distributed interaction.

In Sect 3, we dive deeper into the BBE framework by outlining Randall Beer's adaptive behavior research program (Beer 1997; Beer et al. 1996). This project extends Ashby's insights into a more rigorous methodology that focuses on minimal instances of adaptive behavior, and integrates many of the motivations behind the BBE framework. It does this by combining the insights with modern computer simulation and the mathematical toolset of dynamical systems theory. The goal of Beer's project is to simulate the entire conditions for simple adaptive systems, which includes their environment, their body, and their recurrent dynamical nervous system. He uses evolutionary algorithms to produce dynamical models of brain-body-environment systems that can engage in minimally cognitive tasks, and then analyzes their resulting dynamics to illuminate the dynamical strategies for adaptive behavior. The result of such analysis yields a similar effect as Ashby had intended, of demystifying adaptive behavior by fully reconstructing the system's conditions in a model and then studying its temporal structure.

Most models developed by this methodology have focused on the production of single actions through BBE interactions, and have uncovered the temporal patterns that allow for these particular actions to unfold. But this does not provide a complete picture of living systems, which can generate many possible actions and switch

between them in a context-appropriate manner. Considering the problem of multiple actions brings up new questions about coordination between brain, body, and environment. We ask how multiple actions can arise out of a single absolute system, in which at one time a particular coordination pattern is engaged, and at a different time a completely different coordination pattern is engaged. In Sect 4, we use Beer's methodology to evolve an agent that can generate multiple different actions and smoothly switch between them. An analysis uncovers the strategy that allows it to behave in different ways that requires the coordination of different sensors, effectors, and brain regions.

For the last section we bring together many of the discussed ideas, examine their limitations and suggest improvements for future research. The brain-body-environment framework in cognitive science is a young science and still in its early stages of development. Because of this there are many assumptions left untested and many questions left unexplored. By building up a dialogue and continuing to improve our models, we may someday bring this science from its current emphasis on minimal behavior to the complexities of real living behavior.

10.2 Ashby's Self-Organization in Brain-Body-Environment Systems

In this section we review Ross Ashby's absolute system perspective, from which he believes to have demystified self-organization in his 1962 paper, "Principles of the self-organizing system" (abbreviated PSOS), (Ashby 1962), and with which he presents a scientific framework for the study of adaptive behavior in "Design for a Brain" (abbreviated DFB),(Ashby 1952). The similarity between Ashby's theoretical framework and the one suggested by an integrated BBE framework will be demonstrated by following Ashby's argument as laid out in these two publications, and comparing it to the arguments made by the various fields of research of the BBE. Where Ashby's arguments were based in purely mathematical formalisms, fields within the BBE framework have looked at the structure of particular sensori-motor interactions in the real world and have therefore extended theoretical intuitions into empirical. By following through with Ashby's arguments, we gain a better understanding of how a more integrated BBE framework might someday appear.

Ashby begins PSOS with a definition of a system as an arbitrary assignment of parts, as based on an observer's perspective and not limited to material components. This is a constructivist definition of a system, which is often described synonymously with the term "model." The parts of this system are described mathematically as variables that take a range of states and unfold through time based in the dynamical laws of system. Dynamical rules captures these laws mathematically by defining a transformation rule (evolution operator) that determines how the set of states at one point in time changes to a new set of states in the following point in time. Ashby defines organization as conditionality of variables; "As soon as the relation between two entities A and B becomes conditional on C's value or state then a necessary component of 'organization' is present"(Ashby 1962). An

organized system is one with components whose states are conditional on other states, with dynamical rules that bring this conditionality into effect as the states unfold through time. Central to Ashby's perspective is that the goal of a scientist should be a description of a system as an absolute system, with a mathematical model that is free from internal contradictions. The equations that define a system's states and dynamics needs to be refined and reduced until it is described in a "machine-like way, namely, that its internal state, and the state of its surroundings, defines uniquely the next state it will go to" (Ashby 1962).

With this basis for defining systems and organization, Ashby lays the foundations for the dynamical approach in cognitive science, which is one of the cornerstones of the BBE framework. For Ashby, the ideal description of a system defines its relevant variables, their interdependences, and the dynamical rules that systematically unfold the structure through time. This is the foundational principle of "the dynamical hypothesis" (Van Gelder 1998), which received much attention many decades after Ashby. The dynamical hypothesis proposes a unifying philosophical stance in cognitive science, which insists that cognitive systems are dynamical systems, that they are best understood as dynamical systems, and that therefore scientists should thrive for dynamical explanation of such systems. This stance brings with it a certain understanding; it influences the questions asked, the analyses performed, and the interpretation of results (Beer 2007). Many separate lines of research have come upon this same line of reasoning, and have employed dynamical explanation of cognitive behavior. Dynamical systems have been used to model neuronal system (Izhikevich 2006), entire brain systems (Skarda and Freeman 1987), coordinated motor behavior (Turvey 1990), child development (Thelen and Smith 1996), language (Elman 1995), interaction between language and vision (Spivey et al. 2005), and many more.

As we will see shortly, Ashby rejects the interpretation of self-organization (in PSOS) or adaptive behavior (in DFB) as something that can come out of the internal organization of an organism on its own. He instead attributes it to an opportune matching of organism and environment. The organism alone is a non-autonomous system, whose behavior is partially dependent on its environmental situation. For Ashby, the agent and environment together make an absolute system, and so it is only on this level of description that we can truly understand adaptive behavior. "The organism affects the environment, the environment affects the organism: such a system is said to have 'feedback'" (Ashby 1952). Systems with feedback cannot be treated as if their action was controlled in a linear way; they possess properties that cannot be reproduced in systems that lack feedback. Because of this, if we are to describe an organism's behavior in a model we must bring relevant variables from the environment into the absolute system definition, or else we would miss out on the behavioral effects of feedback.

This essential pairing of organism and environment has been broken up into separate fields in cognitive science and extended in empirical study. The first related field, called situated cognition, finds its roots in the phenomenological philosophy (Heidegger 1962), in classical ethology (Tinbergen 1963), and in ecological psychology (Gibson 1986). Situated cognition concludes that cognition and behavior is a contingent on the situation in which it is enacted (Hutchins 1995; Clark and

Chalmers 1998). A cognitive system is always interacting with its environmental through sensors that perceive, and effectors that produce behavioral output. This leads to the same conclusion as Ashby; behavior is controlled via sensory-motor feedback. Perceptions trigger actions, actions produce changes in the environment, these changes are again perceived and trigger new actions that correct for or extend the effects of the previous actions. Different environmental situations will produce different perceptions, and therefore trigger different actions. This field's emphasis shifts many problems, such as memory and reasoning, from the brain to the environment. Instead of having to conceive, predict, and remember the consequences of an action, action is simply executed by reading off and reacting to information available in the environment.

Embodied cognition is a complementary field, which rather than placing emphasis on the role of the environment, places its emphasis on the role of the body. The significance of embodiment was also first described phenomenological philosophy (Merleau-Ponty 1996), and was recognized by cognitive science in the 1980's with Rodney Brooks' robotics, which emphasized the role of physical embodiment in coordinated behavior (Brooks 1991). Brooks ideas provided a radical alternative to the then-dominant computational approach. Following research in embodied cognitive science has emphasized the role of activity in perception (Noe 2004), autonomy in cognition (Pfeifer and Scheier 2001), the use of metaphor as based in sensori-motor experience (Lakoff and Johnson 1999), and in the philosophy of cognitive science (Clark 1998). These projects recognize how the physical aspects of an organism's body are crucial to its behavior and provide enormous constraints on behavior. The nervous system receives input through the embedding of sensors on its spatially extended body, and their physical properties directly affects that information available to the organism. Additionally, the particular assemblage of bones, joints, and muscles create a unique control problem for the brain. As the cognitive system develops, both physically and behaviorally, it is constrained by a body and can only learn through information provided through a body. Ashby did not place emphasis on the physical embodiment, but would certainly agree with these scientists that the body is a relevant variable, and must be accounted for when describing behavior and cognition from an absolute systems perspective.

The final definition that Ashby offers in PSOS before turning his attention to the demystification of self-organization is what he calls "good organization." This aspect of Ashby's framework has been less influential in the majority of the fields under the BBE umbrella, but has still been advanced in one of the most far-reaching fields called enactive cognition (Varela et al. 1992; Thompson 2007; Di Paolo 2005). For enactive cognition, as well as for Ashby, "good" is a relative term that is dependent on the fit between the system in question and its given environment. If the pairing is such that it acts to further the system's survival, then the system has good organization. Ashby describes what he calls "essential variables" as variables that are closely related to survival (e.g. heartbeat, core body temperature, oxygen level). A successful organism acts to maintain these variables within a narrow range, but when one of the essential variables is significantly altered, the organism dies, and the rest of the essential variables are also dramatically changed. Enactive cognition

introduces a “boundary of viability,” which surrounds the subset of an organism’s state space within which it must remain to survive (called a viable set). It asks how adaptive agents engage with an environment in such a way that they discover possible actions, and engage in actions that bring them to increasingly robust regions of the viable set (Di Paolo 2005).

Now, back to Ashby’s argument in PSOS: with the definitions of absolute system and good organization nicely laid out, Ashby proceeds with his demystification of self-organization. For Ashby, the term self-organization implies a system’s ability for “changing from a bad organization to a good one,” (Ashby 1962). A self-organizing system is one that at first does not have a chance of surviving as an organized system, and by dynamically unfolding through time it changes its organization and is able to persist robustly in the environment. Ashby turns the omniscient lens of the absolute system perspective onto this adaptive behavior, and attempts to describe the dynamical organization that could produce such behavior. He quickly rejects an interpretation of self-organization as something that can come out of a system’s internal organization; “no machine can be self-organizing in this sense.”

For Ashby, organization cannot autonomously improve itself because it is mathematical nonsense to talk of a function that is a function of the state that it defines; an evolution operator, which determines how states unfold through time, cannot be updated by the states that it transforms. There must be some additional variable that drives this apparent organizational change, but which is not contained within the organized system. “The appearance of being ‘self-organizing’ can be given only by a machine being coupled to another machine” (Ashby 1962). Ashby then proceeds to assert that the appearance of self-organization in systems is not only unremarkable in the sense that there are no special conditions that govern self-organization, but that it is in fact an inevitable property of large dynamic systems that have been given sufficient time to come to equilibrium. When we examine this equilibrium, we can split up the relevant components into “organism” and “environment,” and will find that the organism is highly robust to perturbations from the environment, creating what Ashby calls an “adaptive fit.”

This forecasts a demystification of adaptive behavior that would result if there were a complete integration of a BBE framework. By bringing focus to the dynamical approach in which a cognitive system is characterized by a set of states and dynamical operator, proponents of the BBE framework have adopted a fascination with self-organization (Kelso 1995; Maturana and Varela 1980; Thompson 2007). They have recognized that adaptive behavior is the result of feedback between brains, bodies, and environments, and that all relevant variables across this system must be integrated into our models. Just like self-organization cannot come out of a system in isolation, adaptive behavior cannot come out of an organism without an environment to couple with. If we identify the relevant components of a particular cognitive system, and bring them together in a dynamical model, cognition becomes an unfolding process that takes place between distributed components, and can be understood in a purely dynamical terminology. The coupled brain-body subsystem is called the “agent”. It interacts with the environment through coupled interactions that generate feedback. Coupling that flows from the environment to the

agent is called “sensory,” and coupling that flows in the opposite direction is called “motor”. The agent’s behavior is defined by its trajectory of motor outputs (Beer 2007). Cognitive capacities such as memory, learning, attention, and recognition are predicted to fall out of this description if such an ideal model is obtained, not as intrinsic properties of a system but as patterns that emerge from the dynamical trajectories of the system.

10.3 Beer’s Adaptive Behavior Program

Where Ashby had compelling terminology based in dynamics, and a complete vision for the study of adaptive behavior, he lacked on an ability to explain any cognitive behavior that we might find in the real world. Ashby’s example adaptive systems were based in formal systems, which he defined and brought to life by running the equations. He was able to make his conclusions by studying the dynamical properties that arise from such simulation. But the proof that adaptive behavior cannot come out of an isolated system is not sufficient to explain the cognitive behavior we observe in real living organisms. We must ask about their particular structural properties, how these particular properties produce an adaptive fit with the properties of environment, and how the behaviors that we observe result from this opportune matching.

Meanwhile, the BBE framework has approached the cognitive process from an empirical perspective; relevant variables that signify states of real living agents are identified and brought into models that predict how they unfold through time. But the complexity of real living systems is hard to overcome. There are many components, and with their nonlinearity, accurate prediction appears futile. Ashby himself recognized this difficulty, and asserted that real systems likely have infinite variables (Ashby 1952). The difficulty in creating complete models leads to an inability to theorize about adaptive behavior in the terminology of dynamical systems.

A bridge between Ashby’s theory and the BBE’s empirical interest is attempted in Randall Beer’s adaptive behavior project (Beer 1997). What began as the rejection of traditional artificial intelligence led Beer to set as his goal the simulation of an organism’s entire capacities. His research has developed a rigorous methodology to simulate the entire conditions of a brain, body, and environment that engage in minimal instances of adaptive behavior. Beer then adopted dynamical systems theory to analyze the resulting behavior’s dynamical underpinnings. These models are developed by first defining an environment and body in a computer program, and then using genetic algorithms to evolve a dynamical neural network that can control the body effectively in a way that generates the desired action. These agents are of interest to the BBE framework because their simulated behavior is easily related to behavior observed in the real world, yet they are simple enough to be analyzed and completely described with a dynamical terminology.

Beer’s first examples of evolved embodied agents were designed to produce the behaviors of insect walking and chemotaxis (Beer and Gallagher 1992). The chemotaxis agent will be described in greater detail later on in this section. These examples

demonstrated successful situated behavior, but drew some criticism for only studying simple sensory-motor tasks, and not addressing high-order cognitive function. Beer's next step was to extend the framework to simulated examples of minimally cognitive behavior (Beer et al. 1996), which would demonstrate higher-level cognitive behavior as the result of dynamical BBE interactions. Following from this proposal, many agents were evolved to produce behaviors such as selective attention (Slocum et al. 2000), categorical perception (Beer 2003), learning and memory (Izquierdo et al. 2008), relational categorization (Williams et al. 2008a), referential communication (Williams et al. 2008b).

Minimization allows the researcher to focus on a particular cognitive function of interest, and provide this function with a dynamical explanation. By evolving systems that specialize in specific behaviors such as learning, attention, or categorical perception, the cognitive system is reduced to a minimal organization that allows for only the behavior of interest and removes the additional functionality inherent in living systems. This is the concept Beer calls "frictionless brains"; the nervous systems are evolved to produce a well-defined function, which can then be studied without interference by other influencing factors. Real organisms don't have this specialization, and usually take part in many types of behavior whose neuronal underpinning cannot be easily teased apart (this capacity for multiple functions is further addressed in the section on action switching). Because of this, it is much more straightforward to study frictionless brains with well-defined behavioral functions because the neuronal behavior can be directly linked to the production of that particular behavior.

Simple simulated agents are also far more ideal than real living systems for full dynamical analysis. Scientists gain full access to the final, successful agents, because all of the interactions that come into producing the behavior are readily available to the scientist, and just have to be recorded and analyzed. The simulation is a full absolute system model by definition, and does not require a process of abstraction to creating a simplified model. With access to this model, behavior can be analyzed in the way dreamt about by Ashby many decades ago. Beer's framework allows him to ask questions that Ashby could not have begun to answer, such as "How do the individual components across brain, body, and environment contribute to a specific behavior?", "What classes of control mechanisms are best suited to the generation of adaptive behavior?", and "how does manipulation of the variables and parameters affect resulting behavior?"

10.3.1 CTRNNs and Genetic Algorithms

Continuous-time recurrent neural networks (CTRNNs) are adopted as the model nervous systems for these simulations. The general form of these equations is shown below. In this equation, y_i is the state of the neuron, τ is the time constant, w is the weight between neurons j and i , θ is the bias term, I is external input, and $\sigma(x)=1/(1+e^{-x})$ is the standard logistic activation function.

$$\dot{y}_i = \frac{1}{\tau_i}(-y_i + \sum_{j=1}^N w_{ij}\sigma(y_j + \theta_j) + I_i)$$

Beer justifies this selection for a neural model with several points (Beer 1995a). First off, this model's recurrent connections allow the agent to initiate its own behaviors as a result of its internal state unlike the feed forward networks that were popular at the time. Additionally, Beer argues that they are the simplest case of a nonlinear, continuous dynamical neural network model and despite this simplicity, they are universal dynamics approximators (Funahashi and Nakamura 1993). They lend themselves to a biological interpretation, in which a state is associated with a neuron's mean membrane potential, and the output is associated with its average firing frequency. Finally, CTRNNs are computationally and analytically tractable, and they are evolvable by searching through the combinations of the CTRNN's parameter values.

Just like real-world systems produce an adaptive fit between the agent's internal control mechanism and its given environment, so too must the BBE simulation produce a fit between the dynamical nervous system, its body, and its environment. To produce this fit, Beer adopted the use of genetic algorithms (GAs) to evolve CTRNNs that optimize a fitness function by controlling a simulated body in a simulated environment (Beer and Gallagher 1992). This approach was separately developed at around the same time period by some other research groups (Cliff et al. 1993)(Nolfi et al. 2000). These GAs encode the CTRNN parameters, τ , w , and θ in genetic strings. An initial random population of such strings is created, and in each generation the fitness of each individual is evaluated by running a simulation with the individuals' CTRNNs. A new generation is created by selecting highly fit individuals and slightly mutating them to explore nearby regions of parameter space. The selection process chooses individuals with a probability proportional to their fitness, so that more highly fit individuals are represented in the next generation. A set of genetic operators modifies the selected individuals' genetic strings with mutation and crossover. Mutation randomly modifies portions of the strings with some fixed mutation probability, and crossover combines chunks of genetic strings from multiple individuals to create a whole new individual. Once a new population has been constructed, the entire process repeats, and after many such generations the population's fitness increases and converges onto some final local maximum.

The creative part of making these models lies in the experimenter's design of the agent's body and environment, and in the definition of a fitness function that can select for a behavior of interest. If the conditions are designed well, then the GA can move through the space of possible CTRNNs in a gradual way, towards regions of increased adaptive fit. With a good design, a highly successful CTRNN is results from the evolutionary search, and provides the modeler with a BBE system that can be further analyzed to uncover the dynamical strategies that generate successful behavior.

10.3.2 *Dynamical Systems Theory*

The evolved and well-adapted CTRNN is a system of nonlinear differential equations. Each variable makes an axis in the system's state space, which is the set of all possible states that the system can be in. Every point in this state space has an instantaneous trajectory, as determined by the equations, that leads it to a different state (or in the case of equilibrium points, there is no trajectory and the state remain constant). A particular behavior results by setting an initial condition in this space, and following the resulting trajectory through time. A phase portrait is the set of all trajectories that can result in the system. It is the BBE scientist's goal to fully describe this phase portrait, including different factors that shape this space and determine the system's behaviors.

Nonlinear systems such as CTRNNs are difficult, if not impossible, to solve analytically. This makes the characterization of their temporal structure difficult to study in traditional ways. But luckily, the mathematical field of Dynamical Systems Theory (DST) has developed various approaches for characterizing a system's given these constraints. DST has learned that it can uncover much of the systems behavior by focusing on invariant sets within the system's phase space, and on linearized behavior in these invariant sets' direct proximity. Based in this realization, a highly developed set of mathematical tools has been developed, which includes ways to identify a system's invariant sets (e.g. fixed points, limit cycles, chaotic attractors), a characterization of their local structure around these sets (e.g. stability), global structure that connects the sets (e.g. attractor basins, saddle manifolds), and the changes in qualitative structure that occurs with changing parameters (e.g. bifurcations), (Beer 1995b).

Beer adopted the DST toolset to analyze the evolved CTRNNs and uncover how particular trajectories seen in simulated behavior are guided by the system's intrinsic dynamical properties (Beer 1995a). But, as the BBE perspective emphasizes, a characterization of the brain's (or CTRNN's) dynamical landscape is not sufficient to describe the generation of behavior. The CTRNN is not autonomous, and its interactions with the environment cause perturbations of state that would not result from the CTRNN in isolation. Because of this, the CTRNN typically moves through its state space not according to a phase portrait found by a dynamical analysis of the CTRNN alone. This analysis only provides knowledge of the CTRNN's inclinations. For a full explanation of behavior the dynamics of brain, body, and environment have to be brought back together into an absolute system.

By evolving complete dynamical system models of brains, with formally defined bodies and environments, a modeler gains full access to all information about the absolute system and can investigate the underlying dynamical space to explain how a system produces adaptive behavior. Simulated experiments explore multiple scenarios' behavioral trajectories, all of which come together in the absolute phase space uncovered by DST. Each component's contribution can be directly determined with such experimentation: agents can be removed from their environments, and the effects of stimuli on motor output directly analyzed. Connections in the CTRNN can be lesioned, or states held fixed by experimenter control. By doing this, the scientists

can determine how each component contributes to the overall dynamical landscape. Such analyses demonstrate that the CTRNN is attracted to equilibrium points in its state space, and that movement towards these points determines motor output. But as a result of behavioral output, the agent moves within its environment leading to changes in sensory inputs, which in turn alter the CTRNN's phase portrait and influence its next instantaneous behavior.

These ongoing interactions continue to shift the CTRNN's phase portrait either by adjusting the phase portrait slightly, or by creating bifurcations that qualitatively alter the phase space. With these ongoing changes, the system continues to chase a moving equilibrium point. We see that with DST, what were once somewhat vague descriptions of a system's dynamics are given real meaning that can be approached scientifically. Specific actions that are evolved with Beer's method can then be fully described in all their specific nuanced details. By evolving many such agents, each with unique dynamical properties, a general space of strategies is uncovered and builds up a broader picture of adaptive behavior.

10.3.3 A Simple Chemotaxis Agent

In this section we introduce Beer's chemotaxis agents, which were first described in 1992 (Beer and Gallagher 1992), and later extended with a full dynamical systems analysis (Beer 1995a). Chemotaxis is an ideal case of minimal sensori-motor goal-directed behavior, in which agents direct their movement according to chemical signals that are present in the environment. Such behavior is used often by organisms, such as bacteria and nematodes, to approach resources by moving up a chemical gradient. We introduce this basic example now to illustrate an example of the explanation made possible by Beer's framework. We also introduce the chemotaxis agent here because the next section on action switching is based on this early example, but extends it to capture a broader explanation of behavior. By first describing the simpler case, we will be prepared to extend the model later.

For the simple chemotaxis simulations, an agent is given a simple body (shown in Figure 10.1), in this example with a 6-neuron fully interconnected CTRNN. These CTRNN neurons include two spatially extended sensors that detect chemical concentrations at their location, two interneurons, and two motor neurons whose combined outputs produce a torque and thrust, which propel it through the environment. The agent is enclosed in a 2-dimensional square-shaped environment that contains a single circular resource at its center. This resource emits a chemical signal, which diffuses through the environment with intensity proportionate to the inverse square of the distance from the center. The GA's fitness function is to minimize distance between the agent and the resource, which would select for agents that can approach the resource and remain as close as possible for the duration of the simulation.

Beer found that multiple chemotactic strategies evolved under these conditions. We will discuss only the most common strategy here, which is reproduced in a novel agent shown in Figure 10.2. This agent moves forward while turning toward the side on which the chemical signal is stronger. A dynamical analysis demonstrated that




Fig. 10.1 Basic chemotaxis agent morphology, with 6 fully-interconnected CTRNN neurons

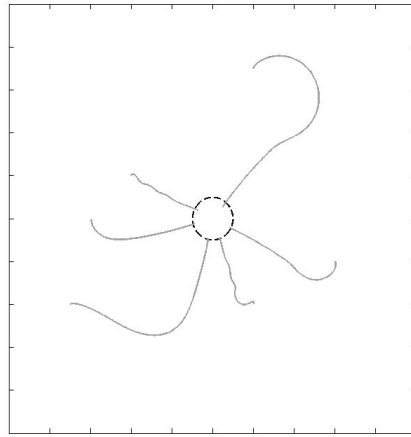


Fig. 10.2 Multiple simulated trajectories, in which the agent is initialized randomly in the environment and moves to the resource

the CTRNN has a single equilibrium point, which shifts with different levels of activation from the two chemosensors. When he examined how this equilibrium point is projected onto the left and right motor outputs, he uncovered a simple explanation for the resulting behavior. When the left chemosensor is more active than the right, the leftward orienting motor neuron becomes more active than the left. When there is higher activation of the right chemosensor, the rightward orienting effector becomes more active. This directly explains the observer turning and approach behaviors, and is rooted in dynamics

10.4 Action Switching

The minimally cognitive behavior project has described many different dynamical strategies for specific cognitive behaviors, but the emphasis on specific behavior

leads to an incomplete picture of organisms' full behavioral capabilities. There are many questions that arise when we shift our focus from the dynamics of performing single actions, to the broad repertoire of actions that all organisms have accessible. Real living organisms depend on the ability flexibly switch between their possible actions. For example, a subtle movement in distant shrubs might be all the information available to a monkey to determine if a predator is on the prowl. This movement in the environment couples with the monkey's relevant sensors, and elicits a dramatic behavioral change from gathering food off of the ground to scampering up a tree for safety. Such actions require very different patterns of sensorimotor coordination; picking food might require fine finger dexterity and acute eye movements for examining food sources. Running up a tree would require full limb coordination and tactile or proprioceptive sensory input. By committing to the dynamical perspective, the scientist is obligated to describe how the many interacting components of the brain, body, and environment become engaged throughout these different actions and in the transition between them.

More traditional approaches in Artificial Intelligence assume that a higher-level mechanism must be used to determine action. This mechanism uses logical or statistical reasoning to decide upon the most beneficial action out of a repertoire of possible actions given the information it has available about the present context. After this decision-making process selects an action, the action is initiated. Some might claim that the BBE approach does contradict this depiction of higher-level mechanisms for action selection. For descriptions that attempt to bring BBE dynamics with higher-level mechanisms of action selection, self-organization and dynamics only account for feedback let loose on a one-way path towards a particular end-state or goal. The initial conditions that allow particular such actions to be instantiated are determined and initialized by a higher-level decision process.

But this misses out on the real underlying message intended by the BBE framework. This perspective aims to describe an absolute system, which has a phase space that describes the systems full range of possible behavior that result from a single, unchanging evolution operator. For this perspective, the agent's behavioral repertoire has to be completely contained within the dynamical explanation, including its movement from the state space region that defines one action to the regions that defines another. Action switching must be a product of self-organization that falls out from these dynamics, and not a higher-level mechanism that sets initial conditions. Many novel questions come up from defining the problem in this way; how is a systems phase space divided between its full repertoire of possible actions. Are there specific regions of this phase space that become responsible for each action (modularity)? How do the different types of action constrain the sensorimotor apparatus in their own unique way that produces appropriate behavior? When an action is completed, how does the system transition to a different action?

Few dynamical systems agents have explored the questions of multiple actions. It has been shown that the same CTRNN is able to implement qualitatively different behaviors when coupled to different bodies (Izquierdo and Buhmann 2008), and that globally stable CTRNNS containing a single basin of attraction are able to sustain multiple modes of behavior (Buckley et al. 2008). In this section we introduce

a new model capable of engaging in multiple different actions, which autonomously switches between these actions without a higher-level mechanism. A highly successful agent is examined in depth to reveal its dynamical organization, and how it allows for efficient action and switching.

10.4.1 *Evolving an Action Switcher*

To study action switching, a chemotaxis agent of the same basic design introduced in the previous section on simple chemotaxis agents was evolved, but with some additions that forced it to switch between the approaches of two different resources. Its environment was encoded as a 100x100 unit plane, with two sources of food that were held in fixed locations for each trial. The agent was given an initial position and directionality in this environment, and was allowed to move spatially by coordinating its motor neurons. Its task was to maintain two nutrient levels above zero by coming within the spatial boundaries of each food source. When it is within the resource’s boundaries, the corresponding nutrient level is increased at a fixed rate to simulate the uptake of nutrients. But there is also a constant decay of nutrient levels at a fixed rate that simulates a simple metabolism. In order to sustain both nutrient levels above empty, the agent must switch between its approaches of each resource, and to spend sufficient time on each one to refill the nutrient supplies.

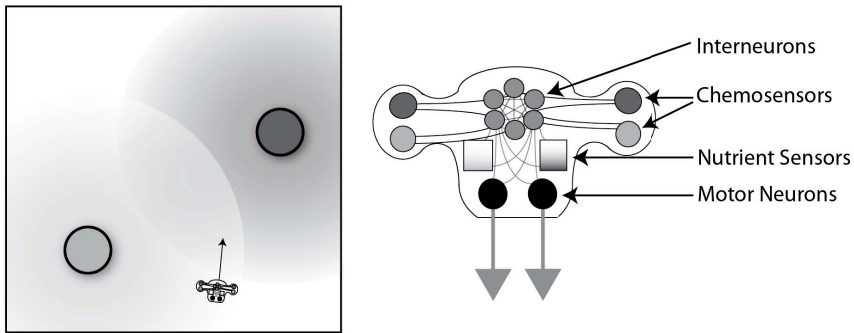


Fig. 10.3 The image on the left shows the action switcher’s environment. The agent exists in a 2-D bounded environment with two different resources and their diffusing chemical signals. The image on the right shows the agent’s morphology. Two types of chemosensors on each spatially-extend stalk detect the different chemicals, nutrient sensors detect internal nutrient levels, and motor neurons produce force that moves the agent.

The agent is given a body that is controlled by the CTRNN via sensors and effectors that were embedded in different locations along its body. Of the neuronal components, four are chemosensors, two are nutrient sensors, two are motor neurons, and six are interneurons. Chemosensors receive input regarding the concentration of a chemical, which just like the original chemotaxis agent, is proportionate to the

inverse square of the distance between sensor and resource. The two chemosensor stalks were extended outward from the body's assigned coordinate location by 6 units, therefor spatially extending the agent's sensors within the environment. The outputs of the two effectors are used to generate movement through the environment. The output of each effector is a vector of force that pushes one side of the body. Directional change is determined by the difference between the two effectors' outputs, while magnitude of movement was determined by the sum of the effectors' outputs. Additionally, the velocity was multiplied by a friction coefficient of 0.9 at every time step, such that the agent would slow down and eventually stop if the effectors produced no force.

The two nutrient levels serve as measures of how much nutrient for each resource is stored in the agent. These values can be anywhere between 0 and 10. When the agent is within the boundary of a resource, its corresponding nutrient is replenished by 0.02 units per time step. Additionally, there is a decay of 0.004 units per time step throughout the entire simulation to simulate a constant metabolism. The simulation continues as long as both internal nutrient values are above empty, but as soon as either value dropped to 0 the agent dies and simulation ends. If the agent is successful and does not die, the simulation is halted after 5000 time steps, which is sufficient time to guarantee that the agent had to move to each resource several times. To further guarantee robustness of the agent in many different environments, it had to succeed in 11 different environments. The average length of time the agent survives in the 11 simulations defines its fitness value for the genetic algorithm.

Figure 10.4 shows the evolutionary progression of twenty different evolutionary runs that had the same morphology but different initial random seeds. The top performing agent's evolutionary run is shown in black. The evolutionary search incrementally increases their fitness as a strategy is converged upon and refined. There are many different possible strategies for action switching given the constraints provided to the GA, and the evolutionary runs can only come upon some of them. In typical experiments of this framework, researchers are recommended to investigate the full space of possible strategies by examining many of the evolved solutions, and also evolving agents with different morphologies (e.g. different amounts of sensors, interneurons, connectivity, etc.). These different morphologies produce constraints on the genetic search, and can lead to very different final strategies. By analyzing multiple strategies, the researcher attains a more general understanding of the behavior, which abstracts over the particular details of an individual instantiation. However, in this chapter only the top-performing agent is selected from the twenty evolutionary runs shown. A more comprehensive examination of different strategies will be described in an article currently in preparation (Agmon and Beer 2013).

10.4.2 The Agent's Behavior and Dynamics

The top-performing agent's final fitness score is .86, which means that it survived through 86 percent of the maximum time provided in the 11 trial configurations. This indicates a high robustness to different environmental conditions, and

persistence in continued action switching between the approaches and eating of the two resources. A typical behavioral trajectory is plotted in Figure 10.5, where the agent moves towards a resource, performs several loops on top of the resource (which we will call eating behavior), then after some time, it leaves the resource and approaches the other resource for a similar eating behavior. We will call the behavior in which the agent moves toward and eats resource 1 action 1, and the behavior towards resource 2 we will call action 2. Accomplishing these actions requires the engagement of different chemosensors that allows for appropriate, directed motor behavior. Somehow, between the two approaches, these different coordination patterns come into effect.

We can attempt to localize action 1 and action 2 by observing and experimenting on the agent’s state during the different actions. This helps identify the relevant vari-




Fig. 10.4 Twenty evolutionary runs, with the top performing evolutionary run in black




Fig. 10.5 A typical behavioral trajectory (solid line) between the two resources (outline by the dashed line)

ables for each action. Part of the trick in this agent's evolution was the introduction of nutrient sensors, which provide input to the interneurons about the two nutrient levels. It is reasonable that the evolutionary search of the GA would take advantage of this information to determine action, and this agent did just that. Figures 10.6 and 10.7 demonstrate how action is dependent on nutrient levels, by showing the behavioral effects of nutrient level manipulation. In Figure 10.6(A), nutrient 1 is held fixed at the near-empty value of 0.5 (out of 10) and nutrient 2 is held full at a value of 10. The agent is shown to approach resource 1, as if it was hungry for the resource, and remains eating the resource for the entire extent of the simulation. Figure 10.6 (B) shows the exact opposite scenario, with nutrient 1 held at 10 and nutrient 2 held at 0.5. We again clearly see that the agent approaches resource 2 and continues to engage in eating behavior for the simulation's full duration.




Fig. 10.6 A) Action 1, the agent moves to resource 1 and engages in eating behavior for the duration of the simulation. B) Action 2, the agent moves to resource 2 and engages in eating behavior for the duration of the simulation.

These demonstrate that the nutrient level plays a role in determining which action the agent is engaged in. But in the reality of a simulation, nutrient levels are constantly increased by eating or decreased by metabolism, such that these idealized actions shown in Figures 10.6 (A) and (B) can never happen. The agent passes through intermediate nutrient states that generate different behavioral tendencies. Figure 10.7(A) shows the behavior when both nutrient levels were held really high, which looks like the agent is not attracted to either resource but instead explores the full environment. Figure 10.7(B) shows behavior when both nutrient sensors are held at a low level, in which the agent moves around in tight circles. These figures also show behaviors that can never be achieved in simulation, because both nutrient levels can never be completely full at the same time, and can't be equally empty. Even though these behaviors are not realistic in a typical environment, they provide interesting insight into the dynamical properties of the system. We gain an understanding of the agent's behavior as ongoing transitions between more explorative modes when nutrient levels are higher, to more exploitative mode when nutrient levels are low. Additionally, off-balanced nutrient levels lead to action directed towards one of the two resources.

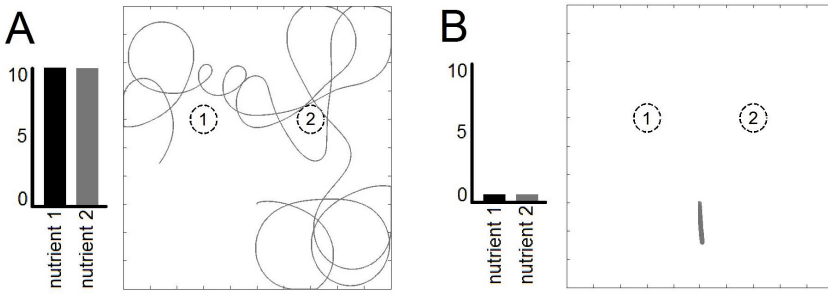


Fig. 10.7 A) The agent explores the environment without approaching either resource. B) The agent moves around in tight circles, while slowly drifting downwards.

Both actions 1 and 2 were further investigated to reveal how sensorimotor interactions are determined by the system, and how this generates the resulting observed behavior. We will focus our explanation on actions 1 by describing Figure 10.8, an identical analysis was done on action 2 and produced similar results, shown in Figure 10.9. We use the constraints of Figure 10.6(A) to approximate the nutrient levels for action 1, with nutrient 1 held at 0.5 and nutrient 2 held at 10. Additionally, it is assumed that the resource 2's chemical trace is irrelevant during action 1, and so the chemosensors for resource 2 are both held fixed at 0. We then calculated the system's equilibrium points as the left and right chemosensors for resource 1 (C_{1L} , and C_{1R}) are varied through their range of states. As it turns out, the system has only one stable equilibrium point for this range chemosensor levels. In Figure 10.8, this equilibrium point is projected onto the left and right motor neurons' outputs (M_L and M_R), which shows the outputs that the motor neurons would tend towards given the chemosensor values were held fixed.

For a reminder of how the total force on the agent is determined: the angular force is determined by the difference between the two outputs. This leads to a slightly counter-intuitive relation between effectors and behavioral output such that when there is more force coming from right motor neuron, the agent turns left. When there is more force coming from the left motor neuron, the agent turns right. When there are equal amounts of force, the agent moves forward with a thrust determined by the effector's added outputs. There is an additional friction constant that reduces the velocity at each time step. With this quick summary, we can interpret the motor surfaces of Figure 10.8. These show that when the right chemosensor is more active than the left chemosensor, the right motor becomes less active than the left motor. This turns the agent to the right. When the left chemosensor is less active than the right, there is tendency to turn left. When both chemosensors are about equally active, both motor neurons are also about equally active, which drives the agent forward. With every movement, the agent's chemosensors move, and elicit new inputs, which through ongoing feedback moves the agent successfully to the resource. These surfaces fully describe the strategy used by the agent, which successfully brings it to resource 1 as shown in Figure 10.6(A).




Fig. 10.8 Individual motor projections of the system’s single equilibrium point for action1 (shown in Figure 10.6(A)). These surfaces are functions of the two chemosensors, $C1_L$ and $C1_R$, which when held at a particular state produce the stable motor outputs in effectors M_L and M_R . Action 2, shown in Figure 10.9, but not identical motor projections.




Fig. 10.9 Individual motor projections of the system’s single equilibrium point for action2 (shown in Figure 10.6(B)). These surfaces are functions of the two chemosensors, $C2_L$ and $C2_R$, which when held at a particular state produce the stable motor outputs in effectors M_L and M_R .

We now turn our attention to Figure 10.10, which examines the agent’s internal states throughout a full simulation of both actions without any imposed constraints. This aims to provide insight into how the internal state behaves in each of the two actions and how it transitions between them. To produce this image, the states of the six interneurons are recorded throughout the same simulation of behavior previously shown in Figure 10.5. There are 6 interneurons, and so their state makes a 6-dimensional state space, which can, for obvious reasons, not be visualized in its completeness here. Instead, a principal component analysis (PCA) is performed on the data, and the top three principal components are identified to make up the axes of Figure 10.5. These components account for 92 percent of the variance of the full interneuron space. The data is transformed into the coordinates defined by these principle components, and projected into the 3-dimensional space. In this plot, the interneurons trajectories throughout two actions are colored in different colors to demonstrate the clear separation of the actions that takes place in the interneuron’s activity. This separation should not be interpreted as a universal rule of action switching, but it is certainly pronounced in this particular case.




Fig. 10.10 This shows the trajectory of interneuron states throughout the two actions shown in Figure 10.5., projected onto a dimensionally reduced space made of the top 3 principle components. The black trajectory is the interneurons' state during action 1 and the gray trajectory is the interneurons' state during action 2.

10.4.3 Discussion of Action Switching

This agent provides a proof of concept that a distributed dynamical system, such as a CTRNN, can control a body in multiple different directed actions and autonomously switch between them. There is no higher-level mechanism imposed, and all the observed actions exist within the same defined absolute system. Though it is a minimal case, it can help set a basis from which dynamical systems terminology can explain more complex examples of action switching.

We gain an understanding of the system as existing in a high-dimensional state space, with specific actions resulting from the agent's state and coupled interactions with an environmental situation. The agent demonstrates that specific actions appear as reductions in the state space, in which some temporarily irrelevant variables can be removed from the analysis. The actions are approximated through projections of the high-dimensional state space onto a reduced subspace made of the temporarily relevant variables, such as $C1_L$, and $C1_R$ left during action 1. But with our analysis of this agent we recognize that these reductions are approximations, and the real system actually exists in the higher-dimensional state space that contains all variables. Action switching is movement through the higher-dimensional space, between these actions' approximated subspaces.

When analyzing how an agent traverses the state space between actions, the modeler needs to demonstrate how a trajectory moves from one projection to another, as when action 1's sensorimotor surfaces, shown in Figure 10.8, would transition to action 2's surfaces, shown in Figure 10.9, which are defined by different variables. This requires movement in more dimensions than we can illustrate directly on paper, and will require different kinds of analyses. In this agent we see the transitions in Figure 10.10, though this image does not illustrate the dynamical landscape, such

as equilibrium points, which would drive the system from one action to another. A more complete description would show the dynamical properties, such as a bifurcation or moving equilibrium point. A more complete picture is provided in a later article (Agmon and Beer 2013).

10.5 The Prospect of Brain-Body-Environment Systems

There is much left to develop and discover in this sapling scientific field; what we have covered in this chapter only provides minimal depictions of what in reality are very complex phenomena. If it is true that we need to take an absolute system perspective by modeling all of the relevant variables underlying adaptive behavior, then the current dynamical system models will need to scale up many orders of magnitude if we are to describe the behavior capacity of living systems. Neuronal, biological, and ecological systems have many relevant variables that come into play in the production of an individual organism's behavior. The actions performed by simulated agents will need to diversify and complexify if they are going to describe these behaviors as seen in reality.

Real repertoires of action are typically much less symmetric than in this paper's example; whereas the agent studied here directed its behaviors towards two resources in functionally very similar ways, real organisms can engage in very qualitatively different types of actions such as reaching, peeling rind off of a fruit, chewing, or fighting an opponent. These actions are coordinated in different environmental situations, by utilizing different sets of sensors and effectors. In dynamical models of such actions, the agent will need to coordinate many degrees of freedom in morphologies that have an increased number of bones, muscles and joints. They will need to gather information about context by integrating different sources of sensory input. Environments will need to expand to include more entities with complex properties that an agent can engage with.

Not only will models of brain-body-environment systems need to scale up to more complex structures and behavior, they will also need to demonstrate adaptivity. Adaptivity is a system's ability to structurally reconfigure itself to behave in increasingly advantageous ways (Di Paolo 2005). As we know from the section on Ashby's self-organization, the change from bad organization to good organization cannot come out of the system's internal organization. Instead we must ask how an adaptive system's robustness increases through experience within the environment. Here, robustness is defined as the maintenance of essential variables (as introduced in the earlier section on Ashby's self-organization). A highly robust system maintains its essential variables far from their boundary of viability, and by doing so reduces its chance of failure or death. When at one time an adaptive system might have fared poorly and closely approached possible failure, at a later time when confronted with the same environmental situation, the system behaves in a more efficient and robust manner.

The model presented in this paper has a static behavioral repertoire, which is given to the agent at its inception by the fixing of dynamical parameters.

Additionally, its body and environment are determined a priori and are unchanged for the duration of simulation. But actions seen in living systems are often acquired through learning and development, during which the environment and body change. Organisms tune their actions and come up with new actions that engage sensors and effectors in novel ways to optimize their interactions. With the absolute systems perspective on BBE systems, learning and development become a type of self-organization that unfolds through interactions across many timescales. Future agents will need to demonstrate this capacity, and their analysis will need to uncover the underlying structures that allow for such abilities.

All of these complexifications will yield new types of behaviors and dynamics that have not yet been described with the methodology introduced here. The field is facing an explanatory gap that it must overcome if it is to provide scientific insight about the adaptive capabilities of life. By starting with minimal models, it has promised to lay a foundation that can be incrementally built up towards more realistic behavior. Minimization has allowed simple behaviors to be analyzed, and their dynamics fully unpacked. But there is no certainty that the types of analyses used in the minimal cases will transfer to more complex instances. The field will need to develop whole new approaches for grappling with these complexifications. It will need new methods for evolving models of behavior that demonstrate diversity, complexity and adaptivity, and new methodologies to analyze their structure and dynamics.

References

- Agmon, E., Beer, R.D.: The dynamics of action switching in an evolved agent. Manuscript submitted for Publication (2013)
- Ashby, W.R.: *Design for a brain*. Wiley (1952)
- Ashby, W.R.: Principles of the self-organizing system. *Principles of Self-organization*, pp. 255–278 (1962)
- Beer, R.D.: Computational and dynamical languages for autonomous agents. In: *Mind as motion: Explorations in the Dynamics of Cognition*, pp. 121–147 (1995a)
- Beer, R.D.: A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* 72(1), 173–215 (1995b)
- Beer, R.D.: The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Systems* 20(2), 257–289 (1997)
- Beer, R.D.: The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11(4), 209–243 (2003)
- Beer, R.D.: Dynamical systems and embedded cognition. In: *The Cambridge Handbook of Artificial Intelligence* (2007) (to appear)
- Beer, R.D., et al.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From Animals to Animats* 4, 421–429 (1996)
- Beer, R.D., Gallagher, J.C.: Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior* 1(1), 91–122 (1992)
- Brooks, R.A.: Intelligence without representation. *Artificial Intelligence* 47(1), 139–159 (1991)

- Buckley, C., Fine, P., Bullock, S., Di Paolo, E.: Monostable controllers for adaptive behaviour. *From Animals to Animats 10*, 103–112 (2008)
- Clancey, W.J.: *Situated cognition: On human knowledge and computer representations*. Cambridge University Press (1997)
- Clark, A.: *Being there: Putting brain, body, and world together again*. MIT Press (1998)
- Clark, A., Chalmers, D.: The extended mind. *Analysis*, 7–19 (1998)
- Cliff, D., Husbands, P., Harvey, I.: Explorations in evolutionary robotics. *Adaptive Behavior* 2(1), 73–110 (1993)
- Di Paolo, E.A.: Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences* 4(4), 429–452 (2005)
- Dupuy, J.P.: *On the origins of cognitive science*. MIT Press (2009)
- Elman, J.L.: Language as a dynamical system. In: *Mind as Motion: Explorations in the Dynamics of Cognition*, pp. 195–223 (1995)
- Funahashi, K.-I., Nakamura, Y.: Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks* 6(6), 801–806 (1993)
- Gibson, J.J.: *The ecological approach to visual perception*. Lawrence Erlbaum (1986)
- Heidegger, M.: *Being and time*. Trans. John Macquarrie and Edward Robinson. Harper, New York (1927, 1962)
- Hutchins, E.: How a cockpit remembers its speeds. *Cognitive Science* 19(3), 265–288 (1995)
- Hutchins, E.: Distributed cognition. In: *Internacional Encyclopedia of the Social and Behavioral Sciences* (2000)
- Izhikevich, E.M.: *Dynamical systems in neuroscience: the geometry of excitability and bursting*. MIT Press (2006)
- Izquierdo, E., Buhmann, T.: Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: walking and chemotaxis. In: *Artificial life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pp. 257–264. MIT Press, Cambridge (2008)
- Izquierdo, E., Harvey, I., Beer, R.D.: Associative learning on a continuum in evolved dynamical neural networks. *Adaptive Behavior* 16(6), 361–384 (2008)
- Kelso, J.S.: *Dynamic patterns: The self-organization of brain and behavior*. MIT Press (1995)
- Lakoff, G.J., Johnson, M.: *Philosophy in the Flesh. The Embodied Mind and Its Challenge to Western Thought*. Basic Books, New York (1999)
- Maturana, H.R., Varela, F.J.: *Autopoiesis and cognition: The realization of the living*, vol. 42. Springer (1980)
- Merleau-Ponty, M.: *Phenomenology of perception*. Motilal Banarsidass Publishe (1996)
- Noè, A.: *Action in perception*. MIT Press (2005)
- Nolfi, S., Floreano, D., Floreano, D.: *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge (2000)
- Pfeifer, R., Scheier, C.: *Understanding intelligence*. MIT Press (2001)
- Polani, D.: Foundations and formalizations of self-organization. *Advances in Applied Self-Organizing Systems*, 19–37 (2008)
- Port, R.F., van Gelder, T.: *Mind as motion: Explorations in the dynamics of cognition*. MIT Press (1995)
- Skarda, C.A., Freeman, W.J.: How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences* 10(2), 161–195 (1987)
- Slocum, A.C., Downey, D.C., Beer, R.D.: Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In: *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*, pp. 430–439 (2000)

- Spivey, M.J., Grosjean, M., Knoblich, G.: Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences of the United States of America* 102(29), 10393–10398 (2005)
- Thelen, E., Smith, L.B.: *A dynamic systems approach to the development of cognition and action*. MIT Press (1996)
- Thompson, E.: *Mind in life: Biology, phenomenology, and the sciences of mind*. Belknap Press (2007)
- Tinbergen, N.: On aims and methods of ethology. *Zeitschrift für Tierpsychologie* 20(4), 410–433 (1963)
- Turvey, M.T.: Coordination. *American Psychologist* 45(8), 938 (1990)
- Van Gelder, T.: The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21(05), 615–628 (1998)
- Varela, F.J., Thompson, E.T., Rosch, E.: *The embodied mind: Cognitive science and human experience*. MIT Press (1992)
- Williams, P.L., Beer, R.D., Gasser, M.: An embodied dynamical approach to relational categorization. In: *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, pp. 223–228 (2008a)
- Williams, P.L., Beer, R.D., Gasser, M.: Evolving referential communication in embodied dynamical agents. In: *ALife XI: Proceedings of the Eleventh International Conference on Artificial Life*, pp. 702–709 (2008b)