

# The Biology of Life and Learning

Christine A. Skarda

James Barham's approach to evolutionary epistemology is written on a grand scale. Nothing wishy-washy about it. Barham goes directly to the big issues: what makes matter alive, what is knowledge, what role does causation play, how can a theory avoid being reductionistic if it subscribes to the ideal of a unified science? I am delighted by the entire project. It got me thinking about the bearing my own work has on some of these issues. The upshot of this interaction follows.

Barham tells us that what biological systems have in common is that they can be modelled as nonlinear oscillators. Consider a simple pendulum. Its dynamic behavior can be modelled in phase space in terms of a periodic attractor: start it moving from any of a number of positions and it will return to the same form of dynamic behavior. This is Barham's "many-to-one" relationship. If input exceeds a certain threshold, e.g., if I reach in and stop the pendulum, oscillatory behavior is replaced with a new dynamic state modelled using a point attractor. The "structural stability" of system dynamics reflected in the ability of the oscillator to converge to the same dynamic state from any number of initial states, and "metastability," which means that the system maintains its dynamic state in the face of perturbations below a certain threshold, are keystones of Barham's theory. Structural stability is said to cause living things to behave purposefully and is the key to a unified science; whereas metastability explains knowledge acquisition, something evolutionary epistemology must address. I'll address the issues of adaptive and goal-directed behavior, as well as learning in what follows.

Barham recognizes the power of the new concepts of nonlinear dynamical systems theory; and he isn't alone, as his extensive bibliography makes clear. The new science of complexity restructures our understanding of nature and of scientific explanation (Skarda & Freeman, 1990). But I believe that Barham draws invalid conclusions from some of the conceptual tools of the new science. I disagree with Barham's claim that the transition from linear to nonlinear system dynamics invests matter with purpose and that structural stability causes goal-directed behavior.

Let's start at the beginning. It is clear that complexity is characteristic of many systems, not just living ones. This is not a new discovery; it was certainly recognized long before Leeuwenhoek produced his marvelous microscope. What was not recognized until relatively recently, however, is that certain complex systems exhibit nonlinear dynamics. This is

---

Christine A. Skarda, Department of Molecular and Cell Biology, University of California, Berkeley, CA 94720

*Journal of Social and Biological Structures* 14(2):221-228.  
ISSN: 0140-1750

Copyright © 1991 by JAI Press, Inc.  
All rights of reproduction in any form reserved.

Barham's point of departure: "the notion of nonlinearity is the key to understanding life." What about nonlinearity? Storms, waves, clouds as well as embryos, slime mold, and brains, to mention just a few, all exhibit nonlinear dynamics. But as this list indicates, nonlinearity does not itself help us grasp what invests matter with life, as Barham claims. The world is full of nonlinear dynamics but only some of the systems exhibiting this behavior are living. Nonlinearity is necessary, but not sufficient to explain life. The good news, as Barham realizes, is that similar forms of system dynamics are characteristic of living and nonliving things. As systems exhibiting nonlinear dynamics, we are just as much a part of nature as are storms and turbulence in the ocean. But this doesn't tell us much.

Obviously the cake has to be cut in a different way if we want to get at the living/nonliving distinction using the tools of nonlinear dynamics as Barham suggests. This brings us to the notion of structural stability. Barham, in a more precise formulation of his position than the one just addressed, tells us that the structural stability of nonlinear systems explains goal-directedness or purpose. This seems a more promising approach. Systems that retain partial autonomy by damping perturbations below some threshold do, in fact, persist over time in spite of the Second Law of Thermodynamics. And the maintenance of structural integrity is essential to biological systems.

But like nonlinearity, structural stability alone cannot explain either the distinction between living and nonliving things or the property of being goal-directed. Take the distinction between living and nonliving systems first. A storm retains structural stability as does the pendulum of my clock, but I do not feel compelled to view these phenomena as living things, nor have many other people for that matter (enthusiasts of thermostats notwithstanding). More is required. As for the claim that structural stability causes goal-directed behavior, a look at the system dynamics of the pendulum certainly refutes this claim. Living systems do exhibit structural stability, but so do storms and pendulums, and we would only refer to these phenomena as "goal-directed" by analogy.

Barham's discussion of structural stability, however, does introduce what I believe to be the deciding factor in drawing the living/nonliving distinction: "self-preservation" or adaptation. Adaptation is not, however, equivalent to structural stability or nonlinearity; nor is it one of the conceptual tools of nonlinear dynamical systems theory. Barham mistakenly conflates adaptivity with structural stability in his discussion of a mathematical description of system dynamics. He identifies the equifinality of trajectories within the basin of a limit cycle attractor as the "self-preserving behavior" of nonlinear oscillators, telling us that nonlinear systems have "preferred solutions." When mathematicians, however, speak of "preferred solutions," they do not mean that the system chooses or "prefers" certain solutions over others. Nor do all nonlinear systems that persist over time do so because they have a global demand for "self-preservation." The pendulum exhibits structural stability under a wide range of perturbations, but if I introduce input to stop it from swinging it stops. What it doesn't do is act to preserve its own dynamic integrity. We require more than the notion of structural stability, more than the notion of nonlinearity, and more than nonlinear oscillators to explain the appearance of structural stability and of goal-directedness.

I believe that Barham is correct to recognize that the bacterium's behavior is unique in that it can only be understood in terms of a global demand for self-preservation. That is, the system must control and maintain itself in optimal condition under the circumstances in which it finds itself. This feature sets apart living and nonliving physical systems whose dynamics may otherwise be identically modelled. My colleague and I have suggested (Skar-

da & Freeman, 1987) that a global demand for self-preservation might function as an internally generated (self-organized) system attractor that sets parameters for all of the subsystems within the organism such that the system produces self-preservative behavior. We call the resulting behavior "adaptive." Adaptive behavior can be distinguished from the reactive behavior of machines and other nonliving physical systems because it introduces a new kind of interaction between the system and its environment. To understand it we need to understand how biological systems function.

Biological systems introduce relations that differ from those of nonliving, physical systems because the conditions with which living things interact as well as their behaviors are determined by an internally generated global constraint for self-preservation. This means that their behavior cannot be explained (like machine behavior or that of the pendulum or storm above) in terms of a response triggered by a stimulus via a feed-forward process of transformation on the input pattern (Skarda, 1986; Skarda & Freeman, 1988)—that model characterizes behavior as a passive reaction triggered by the environment. Because the stimulus has the properties it does, and because it impinges on a system specially tuned to these physical properties, behavior results.

For biological systems, however, interaction with the environment starts within the organism and not with the stimulus. Behavior is determined by internal constraints. For adaptive behavior, the determining factor is the organism's over-arching demand to preserve itself given its unique behavioral repertoire. Consider the bacterium. Climbing up a chemical gradient for the bacterium is not like the behavior of molecules in a liquid. The molecule's movement is caused by the global system dynamics which impinge on it and cause it to behave as it does; the bacterium's movement is determined internally. Both systems are nonlinear, both can exhibit structural stability, but the bacterium exhibits adaptive behavior: it is not simply "moving," it is "acting," and it can do this because its interaction with the world is determined and produced by its internally generated constraint to preserve itself. If, for example, it requires energy, it interacts with the environment as a food source to survive.

An adequate representation of biological reality must tell us not only that the dynamic behavior of the system can be modelled, for example, as an attractor of some sort; it must explain why the organism engaged in one form of behavior rather than in some other form. And this cannot be specified without reference to the internally generated adaptive constraint that determines system dynamics. A complete description of the stimulus and the organism's receptors cannot explain biological interaction. The stimulus cannot be adequately represented merely as a physical structure; it is a biological structure, and its role in behavior can only be understood in terms of the adaptive behavior of the organism as a whole. The same thing can be said for the response. Biological structures exhibit a kind of autonomy that nonbiological ones do not. They not only maintain structural stability, they maintain themselves by engaging in self-preservative, adaptive behavior. The point is this: once the constraint of adaptation is introduced, we are dealing with a special class of physical systems, namely, biological ones. This doesn't mean that some "vital force" has been slipped into the picture. It simply means that biological organization is such that it has evolved to produce true "behavior."

Barham must rework his understanding of the "flow of causal influences" (Figure 2) from stimulus to receptor to effector. The process for living systems starts within the global constraint to survive and with a repertoire of behavioral forms of which the organism is capable. An energy starved bacterium moves through the world defining its environment in

terms of its adaptive behavior as edible or inedible. The last thing, moreover, that the bacterium in search of food needs is to return to the same state. It undergoes a series of state transitions that are continually generated by its global need to survive. For living organisms this means that behavior can be modelled in terms of repeated transitions between unstable states. I suspect that a more realistic picture of biological functioning will focus on instability rather than stability. In my own work I have found the notion of "creative instability" useful. Maintaining a stable state is a good engineering principle because it keeps your bridge or building standing, but I suggest that for organic systems a controlled form of instability is essential. Walking, for example, is really a process of falling forward—an unstable state if there ever was one—that is interrupted briefly when a foot is placed to stop the fall. Stability punctuates this behavior, but does not explain forward movement. Mental states are also examples of creative instability. Stable mental states are characterized generally as pathological, e.g., obsessive/compulsive states.

This brings us to the issue of goal-directedness. Structural stability is no more the property that "causes living things to behave in ways we identify as goal-directed" than nonlinearity "invests matter with purpose," and for many of the same reasons that these concepts fail to explain the appearance of adaptive behavior. Barham has again conflated the issues. Goal-directedness needs to be distinguished from nonlinearity, structural stability, and adaptation. For over a decade I have been involved with the issue of goal-directed behavior. More recently, I have studied the neurophysiology of such behavior in the perceptual system (Freeman & Skarda, 1985; Skarda & Freeman, 1987; Skarda, 1989). These studies indicate that goal-directed activity is unique in several respects. It is essential for learning and memory and other cognitive states, and it is characteristic of organisms that have a unique structural organization and corresponding form of interaction with their environment.

A word of caution before continuing. Certainly Barham's attempt to develop a biologically sound theory of knowledge is laudable. But a unified theory should not ignore differences. "Unified" need not mean "uniform." Biological reality, all physical reality for that matter, is a lesson in diversity. Not only are there all levels of complexity in nature, there are different kinds of organization and interaction. To be "natural," goal-directedness need not be a property of every physical system or even of every biological system. It need only be characteristic of some natural systems. And as I hope to show, it is a feature of physical systems with a very unique structure and form of interaction with the environment. Recognition of diversity is essential for another reason. It prevents us from falling prey to the "Beatrix Potter Syndrome": the attempt to project qualities onto systems at lower levels of organization, and then to see the existence of those qualities as confirmation of the naturalness of the human condition (Rose, 1987). Barham is a victim: "our friend the transpeptidase enzyme" acquires "knowledge" and has "beliefs," a state that Barham says is actually homologous with the states of organisms with central nervous systems and brains. This is not true. Enzymes are part of nature, just as brains are; but enzymes, unlike brains, do not have the "machinery" (pardon the expression) for cognitive states. To understand why, let's examine the neurobiology of learning.

In recent years, researchers have begun to unravel the neural dynamics involved in the goal-directed process of perceptual learning. In our lab, the approach has been to record and measure the neural activity patterns within the olfactory bulb of rabbits before and again after the subject had learned to discriminate two or more sensory stimuli, and to identify the differences in activity patterns that serve to distinguish and classify the neural events with

respect to the discriminanda. Recordings were made of EEG (electroencephalogram) potentials. Analysis showed that odor specific information exists in spatial patterns of amplitude of the oscillatory burst of the respiratory wave (Freeman & Schneider, 1982). Analysis of the EEG traces of the background state before conditioning indicates that every trace has the same temporal waveform, but that the amplitude differs between channels, forming a relatively constant spatial pattern that is unique to a particular animal and that remains constant until odorant conditioning is undertaken. Significantly, no changes in this background pattern occur when unreinforced odorants are presented to animals; however, new patterns do emerge when reinforced odorants are introduced. These patterns can be correlated with particular odors and remain stable within and across sessions provided that the stimulus-response contingencies are not changed. Of particular interest is the fact that these patterns are globally distributed in the olfactory bulb.

The data and the resulting model of olfactory functioning reveal that learning to recognize and recall an odor involves a unique hierarchy of self-organized neural processes that emerge one from the other in a series of state transitions (Viana Di Prisco & Freeman, 1985). This unique process is rooted in what is called the background state. If the subject is motivated (i.e., goal-directed), the following dynamics take place in the bulb. During late exhalation and early inhalation—the period of stimulus input via receptors—intrinsic interaction among bulbar neurons is low. During this stage the activity of afferent (receptor) neurons is imposed on bulbar neurons that accept this information and maintain it by local firing. Learning takes place when a reinforced odorant is presented to the animal over a series of trials, typically a few dozen sniffs. In the bulb, this leads first to the formation of a “nerve cell assembly” (NCA) involving approximately 1% of the neurons in the bulb. After the NCA has formed and as long as the reinforcement contingencies remain unchanged, excitation of any subset of neurons in the NCA by receptors sensitive to a particular odorant will activate the entire assembly. Our model tells us that the background state is a globally distributed, low level chaotic state in which is embedded the local activity of the NCA. (“Chaos” is a term used to describe system dynamics that appear to be random, but are not; chaos is fully deterministic and possesses an order that can be generated by fixed rules that create the complex geometrical forms that represent this form of activity.) We have suggested that the NCA plays a crucial role at the point when receptor input pushes the bulb away from its chaotic rest state to a state change. Using the language of nonlinear dynamical systems theory, we have hypothesized that the NCA determines the new basin of attraction for the system, resulting in stimulus categorization and recognition through the emergence of a globally distributed patterned activity across the entire bulb. These activity patterns, one for each discriminated odor, have been mathematically expressed as a collection of chaotic attractors. These are the patterns sent out of the bulb to the cortex, and that we have suggested are behaviorally relevant for the correlations usually associated with learning and memory (Skarda & Freeman, 1988). Upon exhalation, the bulb returns to its low level chaotic background state in readiness for new, motivated interaction with the environment.

This research tells us that perceptual learning and memory are not passive processes of reaction, like a reflex, in which whatever hits the receptors is registered inside the organism (Skarda, 1989). Perception begins within the organism with internally generated (self-organized) neural activity that, by refference, lays the ground for processing of future receptor input. In the absence of such activity, receptor stimulation does not lead to any observable changes in the cortex. The brain itself creates the conditions for perceptual

processing by generating activity patterns that determine which, if any, receptor activity will be accepted and processed. Perception is a goal-directed dynamic process of interchange inaugurated by the brain, in which the brain opens itself to selected input (and closes itself to others), reorganizes itself, and then interacts to enact change within the organism and/or in the world. I have suggested that the globally distributed, chaotic activity pattern we associate with the stimulus is the one that gives "biological meaning" to input (Skarda & Freeman, 1990).

These data have important implications regarding the physiology of goal-directed behavior. The brains of organisms that exhibit this form of behavior are distinguished by their ability to create a new form of interaction with the environment. In the case of odor recognition, we find that until and unless the animal is "motivated," no state changes occur. Receptors sensitive to an odor may respond, but input to the bulb is ignored. The internally generated state we associate with motivation is responsible for the choice of input and ability to accept it. This is the neural basis of goal-directed behavior.

Not all organisms exhibit purposive behavior. It requires the ability to internally generate the state that, by refference, primes the bulb to receive specific receptor input. This process is accomplished in the "limbic system" of animals from the salamander to the human. The limbic system is a complex system of pathways including the amygdala, entorhinal cortex, fornix, hippocampus and the septum (O'Keefe & Nadel, 1978). Animal experiments have shown that subjects in whom limbic system functioning is disrupted or destroyed exhibit reflex behavior, but are incapable of goal-directed behavior. Why is this?

The limbic system is where the process of goal-directed interaction begins. Here also is the origin of purposive, cognitive states. The process begins with an internally generated motor action pattern which spreads from the brain through the brainstem to muscles; much of it, however, remains within the brain itself in the form of corollary discharge. Some corollary discharge is directed to the sensory subsystems, as we saw in the olfactory bulb, so that they self-organize in expectation of sensory implications of the action about to be taken. The whole system is, thus, united by a single form of behavior in which it is about to engage. This process need not be thought of as involving correlation or comparison, as Barham suggests in his model; rather the entire system can be thought of as moving toward a single form of interaction by way of a process of integration. Past experience is continually reshaped by present experience (learning a new odorant, for example, changes all of the patterns associated with previously learned odorants). Integration of this in the limbic system assures the system's unified action in the present. Interaction, of course, will mean something different for each subsystem. Notice that the organism so alerted is not simply breathing or inhaling, it is sniffing and searching. Perception, at the neural level, is a global burst of internally generated, patterned neural activity that has been chosen and then shaped by the corollary discharge and interaction within the limbic system (Freeman, 1990).

The goal-directedness of this process is not observer relative, as was the case with Barham's enzyme *T*. This behavior actually is goal-directed because it involves a special form of organizational dynamics, something only found in the limbic system of a subgroup of biological systems. Enzymes don't have the necessary machinery—any more than do individual neurons within the brain. It follows that goal-directedness is also different from adaptive behavior. All living things evolved and show adaptive behavior, but not all of these have evolved with the brains and limbic systems required for goal-directed behavior. On the other hand, goal-directed behavior is not always adaptive. We need only look at our own

species for clear examples of nonadaptive, goal-directed acts. And finally, it goes without saying that humans are not alone in their ability to behave with a purpose. The salamander and the octopus, as well as our dogs and songbirds, all enjoy the same privilege.

One more point before concluding. Barham cautions that his use of the concepts of nonlinear dynamics are for "heuristic purposes only" and that "no claims to scientific rigor are intended." But the notions do play an explanatory role in his account of knowledge, and as such I have tried to examine their implications. I have not analyzed his use of these concepts per se. I would however, like to make a few observations in closing. Not all system dynamics are oscillatory. There are different kinds of oscillatory behaviors, but there are systems as well whose dynamics are radically different and which in no way represent a "sub-category" of oscillatory behavior as Barham suggests (note 5)—systems, for example, that can be modelled using point attractors and chaotic attractors. Brains exhibit the latter form of dynamic behavior. As mentioned above, chaos plays a central role in perceptual learning and memory. But chaos is not simply an inevitable consequence of a highly interconnected, complex system. Chaotic dynamics achieve something that Barham's oscillatory dynamics cannot.

Mechanical information processing systems use periodic or steady state dynamics like those in Barham's model. These systems require filters designed by their creator to define in advance what is signal and what is noise. Without this design feature they cannot process information because the "information" has not been defined for them. Brains have to accomplish this task themselves in the face of infinite environmental complexity, and the definition of what constitutes a signal cannot be imposed from the outside. The role of chaotic neural dynamics is to create information, to separate on any given occasion what is signal and what can be ignored as noise. In this way, chaos sets brains apart from all other information processing systems (Skarda & Freeman, 1990). As we saw in the olfactory bulb, the background state that prepares the sensory subsystem for perceptual interaction, as well as the globally distributed activity patterns that we associate with particular odorants, are chaotic states. Chaos is ubiquitous in the process of perceptual learning and memory, as well as in the processes leading to perceptual interaction with the environment. Given the evolutionary epistemologist's concern with cognitive states, this information is highly relevant. As Barham says, "living things are not mechanisms." Neo-Darwinism did not get at the problem of life and evolution, Barham claims, because it remained wedded to a mechanistic biology. But periodic oscillatory behavior is not the way to escape this dead end, since this is just the kind of behavior that characterizes machines. Biological systems that have knowledge not only possess limbic systems with densely interconnected feedback pathways and internally generated global dynamics, they use chaotic dynamics in the creation of knowledge.

To conclude, Barham suggests two criteria for judging his theory: the first of these being that it must "provide a scientifically-satisfying account of knowledge as a natural phenomenon." I think Barham's account falls short of meeting this standard for one reason: it ignores biology. Barham employs the fascinating concepts of nonlinear dynamical systems theory and has the necessary philosophical background to deal with the epistemological issues, but his evolutionary epistemological approach to knowledge as a natural phenomenon requires, first and foremost, an understanding of its biological basis. After all, knowledge and other cognitive states are characteristic of certain biological systems. Barham chooses instead to base his interpretation of knowledge on the "epistemic engine." In other words, he adopts a machine model. This substitution, unfortunately all too common in philosophical circles

(such as Churchland, 1986), is illegitimate for the reasons cited above. Machines just do not have the right form of interaction with their environment, nor the right sort of system dynamics. Neuroscientific research reveals the mistaken assumptions involved in such an approach. Evolutionary epistemology requires the tools provided by neuroscience in order to draw biologically accurate conclusions about the nature of knowledge and goal-directed behavior. Only with these tools can we begin to address the real issues raised by evolutionary epistemology and the philosophy of mind.