Chaos and the new science of the brain. *Chapter in Shaw G (ed.) Concepts in Neuroscience 2 275-285. Skarda CA, Freeman WJ (1990)*

CHAOS AND THE NEW SCIENCE OF THE BRAIN

CHRISTINE A. SKARDA and WALTER J. FREEMAN

Concepts in Neuroscience, Vol. 1, No. 2 (1990) 275–285 World Scientific Publishing Company. Received March 16, 1990

Keywords: Biophilosophy, brain theory, chaotic dynamics, cooperative neural mass, neural networks, neurophilosophy, oscillations, perception.

ABSTRACT

Neuroscience involves hard work, but it is also a lot of fun, especially when it gets its hands on ideas that explain old facts in surprising new ways [1]. This is what has happened during the last decade. With the introduction of nonlinear dynamical systems theory – especially the theories of self–organization and chaos – neuroscience has acquired powerful new concepts for analyzing and interpreting data. In our laboratory this has resulted in a radical transformation of our understanding of cortical functioning, but the application of nonlinear dynamical systems theory in neuroscience has further implications: it forces a revolution in the practice of neuroscience. In our research, the realization that self–organized and chaotic dynamics are essential to brain function has led us to reject the underlying explanatory framework that made reductionism the hallmark of scientific explanation. What is emerging today is not only a new view of brain function, but a new science of the brain.

1. Alternative Approaches in Neuroscience

To understand the revolutionary impact that the theories of self–organization and chaos have had on our model of cortical functioning and on our understanding of neuroscientific explanation, we need to review briefly alternative approaches. The brain is a physiochemical system that operates simultaneously at many hierarchical levels. Neuroscientists seeking to discover the physiological basis of behavior have located it at various levels of this hierarchy.

Until recently, many if not most neuroscientists believed that the physiological basis of behavior was to be found at the level of individual neurons. This view, known as the neuron doctrine [2] or single unit approach to brain functioning, assumed that behavior could be explained in terms of the activity of individual cells triggered by a stimulus. Researchers, recording trains of action potentials from single cells, found that certain cells, "feature detectors" (2], were maximally driven by specific stimuli. Further research showed that these cells were organized in geometrical arrangements that account for their functional properties, e.g. the neuronal "map" of the retina in the visual cortex. The resulting task for single unit theorists has been to determine the patterns of forward connections among cortical neurons that can explain how complex features such as line segments or angles are synthesized from simple features such as on–center off–surround cells [3].

Other researchers believed that the explanation of behavior should be sought in phenomena at much lower levels of description. For example, researchers such as Hyden [4], Rall and Rinzel [5], and Lynch [6] suggested that the lasting changes responsible for learned recognition and memory occurred at the level of the neuronal organelle, such as the synapse or dendritic spine. Still other researchers have designated biochemical changes at the synapse as the biological basis of behavior, extending the search within the neuron on the premise that any permanent change that subserves learning may involve changes in the genome or a modification of the RNA in the Nissl substance (4,7).

A different approach has recently been taken by researchers who postulate that the changes involved in learned behaviors, although based on or involving cellular and molecular modifications, are widely distributed spatially and should be understood first at the level of the neural network [8–14]. Neural network models identify the biological basis of behavior with a distributed process that takes place by gradually varying connection strengths among units comprising the network. Moreover, according to at least one school of neural network theorists, the biological basis of behavior is not only globally distributed in the network, it is a self–organized process that requires the use of the analytic tools of nonlinear dynamics, unlike the approaches mentioned above. Many exciting discoveries relating to the formation and functioning of neural networks have resulted from this approach in recent years, and it has encouraged a very fruitful interchange of ideas among neuroscientists, physicists, mathematicians, engineers, and cognitive scientists.

Despite their obvious differences, these approaches have several features in common that set them apart from the view of cortical functioning developed in our laboratory [15]. First, they assume that the biological basis of behavior can be explained in terms of the properties of the system's parts, whether these be connection strengths among units in the network, individual neurons, or at the level of the genome. Even neural net models that recognize the role of self–organized dynamics focus exclusively on the mechanisms of synaptic change in network components to the exclusion of other possible mechanisms, mechanisms that do not involve synaptic or other changes at the level of the system's parts.

Second, these approaches view neural functioning as a passive reaction to the stimulus. Just as feature detectors react to the stimulus that drives them maximally, so too the models developed by neural network researchers process whatever information is received by receptor level neurons. The resulting activity may or may not be understood as involving self–organized dynamics. None of these models view brain functioning as active or selective. For these approaches, behavior is understood on the reflex model of physiological functioning as a reaction to stimulus input.

Data from our laboratory tell us that in these and other respects, the alternative approaches do not adequately characterize neural functioning [16–18]. We have found that brain function cannot be explained in terms of features of neurons taken individually or as part of a local network, nor is it adequately characterized as a passive reaction to stimuli. And while neural network theorists use nonlinear dynamics in modeling their networks and recognize that self–organization plays a role in the brain, they have yet to realize the radical implications of the concept of self–organization both with respect to their explanatory models and for the practice of neuroscience. The following sections describe the neural dynamics of our model, the functions served by chaotic dynamics in the brain, and the implications of our findings for neuroscience.

2. Chaotic Dynamics in Cooperative Neural Masses

Faced, like other researchers, with the hierarchical structure of neural functioning, our approach has been to investigate neural dynamics at that level of organization in the hierarchy that corresponds in its time and distance scales to the coordinate systems of the behavior studied. In our research on pre–attentive perceptual recognition and memory, we know from measurement of response latency that it takes place within a few tenths of a second after sensory input is transmitted to the cortex. Lesion studies tell us that the neuronal correlate of this behavior is activity in large cortical areas, with a time scale of a fraction of a second, expressed in spatially extended patterns of activity.

We have coined the term "cooperative neural mass" to express this level of neuronal functioning [19]. It is largely thanks to the analytical tools of nonlinear dynamics that we have been able to measure and interpret these spatially extended patterns of activity in the nervous system. Our approach has been to record and measure the neural activity patterns within the olfactory bulb before and again after a subject had learned to discriminate two or more sensory stimuli, and to identify precisely the differences in activity patterns that serve to distinguish and classify the neural events with respect to the discriminanda.

In the experiments [20], thirsty rabbits were conditioned to lick in response to an odorant followed after two seconds by delivery of water, and just to sniff in response to an unreinforced odorant. Recordings were made of EEG (electroencephalogram) potentials using a chronically implanted 8X8 array of electrodes (spacing: 0.5 mm) covering approximately 20% of the surface of the olfactory bulb. The typical pattern of the bulbar EEG was a slow wave, called a respiratory wave, with a burst of oscillation in the gamma range (35–90 Hz) common to all 64 channels. Analysis revealed that odor specific information existed in spatial patterns of amplitude of the oscillatory burst. Analysis of the EEG traces showed that in the background before conditioning, every trace had the same temporal waveform, but that the amplitude differed between channels forming a relatively constant spatial pattern that could be easily identified with a particular animal and that remained constant until odorant conditioning was undertaken. No changes in this background pattern occurred when unreinforced odorants were presented to the animal; however, new patterns did emerge with reinforced odorants. These patterns remained stable within and across sessions provided the stimulus–response contingencies were not changed. Of particular interest is the fact that these patterns were globally distributed in the bulb. The information that served to classify them correctly could not be localized to a particular subset of channels [22].

Our job was to produce a biologically sound model of the background state and of the emergence of globally distributed, odor specific spatial patterns. The resulting model is derived from studies of changes in the waveform of these evoked potentials, and on their replication by nonlinear differential equations simulating. the dynamics of the bulb, anterior olfactory nucleus and prepyriform cortex [23].

The data and the resulting model of olfactory functioning reveal that odor recognition and recall involve a hierarchy of self–organized neural processes that emerge one from the other in a series of state transitions. The hierarchy is rooted in what we have termed the background state. 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 neurons is imposed on bulbar neurons that accept this information and maintain it by local firing. Learning occurs when a reinforced odorant is presented to the animal over a series of trials, typically a few dozen sniffs. In the bulb, this involves first of all the formation of a nerve cell assembly (NCA). The model tells us that excitatory neurons, synaptically linked by bidirectional synapses, become coactivated in pairs upon presentation of a reinforced stimulus strengthening their joint synapses in accordance with Hebb's rule [9]. This leads to the formation of a NCA for a particular odorant that involves about 1% of neurons in the olfactory bulb. After the NCA has formed and so long as the reinforcement contingencies remain unchanged, excitation of any subset of neurons in the network by receptors sensitive to a particular odorant will activate the entire assembly. Our model tells us that this background state is a low level chaotic state in which is embedded the locally disseminated activity pattern of the NCA [24].

We have suggested that the NCA plays a crucial role at the point when receptor input pushes the bulb away from its rest state to a state change. We see its role as threefold: (1) to accomplish the difficult task of generalization over equivalent receptors, to amplify and stereotype the small input received on any given inhalation; (2) to produce the locally disseminated, low density activity pattern in the NCA upon interaction with a stimulus; and finally (3) to provide the mediating mechanism upon state change for the emergence of the globally distributed, odor specific activity pattern we associate with a particular odorant. Using the language of nonlinear dynamics we have hypothesized that the NCA determines the

Our model explains the emergence of these globally distributed activity patterns in the following manner. Receptor input to the bulb does more than facilitate the formation of the NCA. During late inhalation, input to the bulb not only activates the subset of neurons involved in the NCA, it excites all bulbar neurons increasing their strength of interaction, priming the entire bulb for an explosive and sudden state change. Receptor input, thus, destabilizes the bulb; it augments interaction over the entire bulb by pushing bulbar neurons far from their initial low energy state. The result is a state change or bifurcation that leads to the emergence of a globally distributed, odor specific activity pattern. Upon bifurcation, the bulb converts to a transmitting mode in which bulbar neurons no longer respond to receptor input. In this state, information carried by each neuron is disseminated over the entire bulb and integrated by every neuron in the bulb. These patterns of globally distributed activity, one for each discriminated odor, have been mathematically expressed as a collection of chaotic attractors. These are the patterns that are sent out of the bulb to the cortex and that we suggest are behaviorally relevant for the correlations that are usually associated with learning and memory. Upon exhalation, the bulb returns to its low level chaotic background state in readiness for new interaction with the environment [25].

3. The New View of Perception

What are the implications of this model for our understanding of the nature of perceptual processing in the brain? We believe that they axe far–reaching and seriously undermine alternative models of cortical functioning [17,18]. Once it is admitted that perceptual processing involves self–organized, internally generated neural processes, we believe that the classical model of physiological functioning must be jettisoned. The idea that perception can be explained in terms of feedforward processing, that it is caused by the stimulus or can be explained as the sum of responses to stimuli, is no longer acceptable [16]. Our model tells us that perceptual processing is not a passive process of reaction, like a reflex, in which whatever hits the receptors is registered inside the brain. Perception does not begin with causal impact on receptors; it begins within the organism with internally generated (self–organized) neural activity that, by re–afference, 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. It is the brain itself that creates the conditions for perceptual processing by generating activity patterns that determine what receptor activity will be accepted and processed. Perception is a self–organized dynamic process of interchange inaugurated by the brain in which the brain fails to respond to irrelevant input, opens itself to the input it accepts, reorganizes itself, and then reaches out to change its input. We suggest that the self–organized process that replaces environmental input with an internally generated, chaotic activity pattern is one that gives "biological meaning" to the stimulus.

> Perception does not just "copy" objects, it creates their meaning for the organism: "(the) function of the organism in receiving stimuli is, so to speak, to 'conceive' a certain form of excitation" [26].

Our model tells us that the globally distributed activity patterns we record in the olfactory bulb are the neural basis of biological interaction: what happens in the brain is about interaction. Motivation involves the creation of a self–organized internal state that destabilizes the system so that it becomes ready to respond to a specific class of stimulus input within a given sensory modality. This class of stimuli may be quite general and may or may not have been experienced before, but once an exemplar is received it sets up conditions such that the system win generate new forms of interactive behavior to cope with the constraints imposed by new circumstances and previous experiences. Perception is an interactive process of destabilization and re–stabilization via self–organized dynamics [18]. Thus, we come to view the brain as the location where a self–organized process of patterning takes place, a process that reaches back toward the stimuli giving them form at the same time as it creates their biological meaning for the organism.

4. The Contributions of Chaos

It was once generally assumed that chaos was undesirable, that it occurred in brains subject to pathological malfunction, and that 'normal' physiological functioning resulted from dynamic processes that could be modeled as periodic. Our data suggest the opposite view: deterministic chaos is essential to normal brain functioning at many levels of activity. What we previously dismissed as "noise" in the system, something to be eliminated with filters when recording, something that the brain seemed to be fighting an impossible battle with in information processing, now appears to be the behaviorally relevant signal [17, 23].

Having discovered that chaotic activity is ubiquitous in neural functioning, we have asked ourselves: what is it doing? What advantages does chaotic activity confer on brains interacting with their environment? In other words, why chaos? What can it do that other forms of dynamic activity cannot?

We have postulated several important benefits of chaotic activity [17]. One class of benefits concerns the system's biological functioning. It is a fact about brains that their neurons must be exercised in order to assure their proper functioning or they die. We have suggested that the chaotic basal activity of the background state provides a suitable biological mechanism for this; moreover, one that is reliable because it is independent of stimulus input. The brain is built to ensure its own steady and controllable source of noise that is quite stable, but not absolutely so. We have also suggested that chaotic activity enables the rapid state transitions essential for information processing. Without this ability, the brain could not quickly concern itself with a new task. Thus, we can thank chaos for the rapid transitions between perceptual states. Without it, perception would be agonizingly slow. We have also suggested that chaos is the mechanism whereby potentially fatal, and hence undesirable, periodic cortical behaviors are desynchronized: "(if) one wanted to desynchronize a process, the availability of a chaotic attractor would offer an opportunity to do it by a low–dimensional control" [27].

A second class of contributions concerns the ability of brains to generate information. Chaos has a role to play that sets brains apart from all other information processing systems. Chaos is not just an inevitable consequence of a highly interconnected complex system, it is essential for the creation of information. The brain, unlike machine systems, is selective, i.e., it does not process whatever information is received at the receptor level. As we have seen in the olfactory system, unreinforced odorants do not cause neural activity in the bulb: receptor level activity only leads to the formation of a NCA and bifurcation to the global activity pattern when the stimuli are reinforced and the animal is "motivated". This selection of relevant information is not imposed on the system from the outside, as is the case in machine systems which use periodic or steady state dynamics and require filters designed by their creator to define in advance what is signal and what is noise. Brains have to accomplish this task themselves in the face of infinite environmental complexity.

Our model suggests that selection results from chaotic bifurcation. As we have described, a self–organized chaotic generator responds to environmental input by replacing it with an internally generated chaotic activity pattern. These self–organized chaotic activity patterns are transmitted further into the brain and provide the basis for future selectivity by (1) causing changes that mediate motivation, reinforcement and learning, and (2) modifying receptor input by causing direct environmental manipulation by the organism or by changing receptor positioning with respect to the world. The brain determines which input it will admit and what spatiotemporal form the resulting neural activity will assume. We suggest, therefore, that chaos is essential for input selection, processing and the creation of information in the brain.

The interplay of chaotic dynamics among neural subsystems allows the brain to do what no man–made system has yet remotely approximated. It is this hierarchically arranged interplay of internally generated, chaotic dynamic activity that puts the neural information processing system in a class by itself.

5. Neuroscience in Transition

As mentioned earlier the practice of neuroscience, not only its content, must undergo a profound transformation as a result of the introduction of the analytical tools of nonlinear dynamics. Before concluding, let us take a closer look at several aspects of this transformation.

Our data and the resulting model tell us that brains use chaotic dynamics. This finding has implications with respect to the methods of data analysis used in neuroscience. In the past, when the accepted view of neural functioning assume that the behaviorally significant neural events could be understood as periodic or steady state phenomena, researchers relied heavily on Fourier analyses, Wiene and Kalman filters, and autoregressive analyses when modeling and analyzing their data. But once the essential nonlinearity and chaotic character of neural activity is accepted, these analytical methods are no longer adequate. Researchers must adopt new methods, such as reconstructing attractors, in order to understand system dynamics that cannot be accessed by previous methods of analysis.

Second, recognition of the essential role of self–organization in brain dynamics brings with it the need to adopt an explanatory framework that is alien to that traditionally used in science [16]. Self–organizing phenomena, such as fluid dynamics and embryonic development, traditionally resisted attempts to explain them in reductionistic terms, i.e., to explain system properties, like turbulence, in terms of the properties of parts of the system.

It was assumed that the elements of explanation must mirror the compositional structure of the system. Reductionism could not accept that phenomena are simultaneously individual and part Of a greater whole; it claimed that ideally explanations of higher order phenomena would be collapsed into lower order ones and that lower order phenomena were the ultimate explanatory elements, the "causes" that science sought to isolate. Yet the breakthrough of nonlinear dynamics has shown us that explanations of self–organizing phenomena can only be given in terms of the qualitative forms of behavior of the system as a whole, i.e., in terms of system properties that resist analysis in terms of the properties of the parts, whether they be individual neurons or discrete input to the system. This implies that in explanations of self–organizing brain dynamics, there necessarily will be relative independence from the nature and properties of the substrate; hence micro–reduction, the aim of traditional explanations, does not work [28].

The observation that brains employ not only self–organization but chaotic dynamics to produce behavior places yet another nail into the coffin of reductionism. Chaotic phenomena preclude long–term predictions. It may seem paradoxical that a deterministic phenomenon is inherently unpredictable, but in systems that exhibit chaotic behavior, small uncertainties are amplified over time by the nonlinear interaction of a few elements. The upshot is that behavior that is predictable in the short run becomes intrinsically unpredictable in the long term [29]. As a result, physiologists cannot make strict causal inferences from the level of individual neurons to that of neural mass actions, nor from the level of receptor activity to internal dynamics. The causal connection between past and future is cut.

One final point: The rejection of the reductionistic explanatory paradigm also has implications for our understanding of the relationship among neuroscientific research undertaken at different levels of description. Our model suggests that information processing subserving learning and memory involves single units, cell assemblies and mass action among large populations located in neural subsystems. With respect to experimental approaches, this implies that various levels of research will be required to give a complete explanation of the biological basis of learning and memory. Biochemical or single unit studies cannot explain or address the coordinated and distributed changes in large populations of neurons or in the NCA, and EEG studies of global activity patterns will not explain the synaptic changes leading to the formation of the NCA. An adequate understanding of phenomena at a particular level is only obtainable given the methods and concepts proper to that level, e.g. certain aspects of an action potential are only understandable in physiological but not biochemical terms.

However, a clear explanatory hierarchy does exist in neuroscience. For example, while it is true that all behaviors can be mapped onto biochemical and physiological changes, not all such mappings will be useful or even relevant as far as an explanation is concerned: not all biochemistry maps onto behaviors of the individual organism [30]. Thus, we can ascribe an order to our investigations, but it is essential not to fall into the reductionistic interpretation of it outlined above. Because a level is lower in the explanatory hierarchy does not mean that it "causes" the higher order phenomena. The explanatory hierarchy reflects, in part, the relative independence of system properties from lower order phenomena, such as the emergence of global activity patterns that do not themselves involve further synaptic changes like those required at a lower level for the formation of the NCA. Thus, when researchers study a learned behavior such as odor recognition, biochemists, physiologists and cognitive scientists all study the *same phenomenon,* but each studies it at a different level of description. It is a mischaracterization to understand this hierarchy in causal terms. Contrary to popular opinion, feature detectors do not 'cause' perception of objects any more than neurotransmitter imbalance 'causes' mental illness. These forms of causal thinking rely on a reductionistic understanding of scientific explanation that is no longer tenable.

6. From Biophilosophy to a New Neuroscience

Nearly half a century ago, the biophilosopher Merleau–Ponty [31] proposed an alternative explanatory framework for physiology that was truly revolutionary; so revolutionary, in fact, that fifty years later it stands head and shoulders above contemporary work in the field of "neurophilosophy" [321. After examining the neuroscientific research of his day, Merleau–Ponty concluded that physiologists had systematically misrepresented brain function because they were wedded to an explanatory framework that distorted their findings. He argued for a new view of cortical functioning. He claimed that the merely transformational and reactive processes that had been isolated by researchers in the artificial setting of the laboratory, in no way proved that such processes operated in the intact, freely behaving animal. This passive, reflex–based view of physiological functioning was, he claimed, an illusion created by physiologists who tried to understand the brain as a mechanical device. He proposed that brain function should be understood instead as basically creative and selective, and suggested that behavior was inaugurated within the organism rather than by the stimulus. He also argued against reductionistic explanations of neural functioning that reduced system properties to the sum of the properties of its parts, and that understood the components as the underlying causes of the behavior. He suggested that there existed internally generated, global states of cortical activity that could not be explained in reductionistic terms. He referred to these states as 'holistic'; and while his conceptual arguments for them were convincing, it was never clear what the physiological correlates of such states were or how they were internally generated. Revolutionary as Merleau–Ponty's theories about the brain and neuroscience were, they remained just theories, rejected by scientists and philosophers alike as "unscientific".

Today, thanks to the advent of nonlinear dynamical systems theory and its methods of analysis in neuroscience, we have the conceptual apparatus to grasp the truth of Merleau–Ponty's biophilosophy in terms of a new neuroscience. Today, we have at our disposal the tools to discover and explain the internally generated, self. organized, distributed phenomena that Merleau–Ponty could only hint at; and we have developed recording techniques to access these phenomena in intact animals. In addition, the theories of self–organization and chaos have given us a nonreductionistic explanatory framework for understanding brain function. Consequently, we can now explain why reductionistic thinking cannot adequately represent the physiological reality of cortical functioning. And finally, having rejected reductionism, we are in a position to reconceive the relationship among neuroscientific investigations undertaken at various levels of description. Thus, a revolution is taking place in neuroscience today, a revolution that, as Merleau–Ponty suggested half a century ago, promises to completely transform our understanding of brain function and the structure of scientific explanations.

References

[1] S. J. Gould, *Hen's Teeth and Horse's Toes: Further Reflections in Natural History (*W. Norton & Company, New York, 1983). [2] H. Barlow, *Perception* 1, 371 (1972).

[3] D. Hubel and T. Wiesel, J. *Physiology* 160, 106 (1962). [4] H. Hyden, *The Neuron* (Elsevier, Amsterdam, 1967)_

[5] W. Rall and J. Rinzel, *Biophys.* J. 13, 648 (1973).

[6] G. Lynch, *Synapses, Circuits, and the Beginnings of Memory* (MIT Press, Cambridge, Massachusetts, 1986).

[7] E. Kandel and J. Schwartz, *Science* 218, 433 (1982).

[8] W. McCulloch and W. Pitts, *Bull. Math. Biophys.* 5, 115 (1943).

[9] D. Hebb, *The Organization of Behavior* (Wiley, New York, 1949).

[10] F. Rosenblatt, *Principles of Neurodynamics,* Perceptrons and *the Theory of Brain Mechanisms* (Spartan, New York, 1962).

[11] S. Grossberg, *Mathematical Psychology and Psychophysiology* (American Mathematical Society, 1981)

[12] S. Amari, *IEEE Trans Syst. Man Cybern.* SMC–13(5), 741 (1983).

[13] T. Kohonen, *Self– Organization and Associative Memory* (Springer, Berlin, 1984).

[14] J. Hopfield and D. Tank, *Science* 233, 625 (1986).

[15] C. Skarda and W. Freeman, in *Systems with Learning and Memory Abilities,* ed. by J. Delacour and J. Levy (Elsevier, Amsterdam, 1988) P. 199.

[16] W. Freeman, *Perspect. Biol. Med.* 24, 561 (1981).

[17] C. Skarda and W. Freeman, *Behav. Brain Sci.* 10, 161 (1987).

[18] C. Skarda, *La Nuova Critica* 9/10, 49 (1989).

[19] W. Freeman, *Mass Action in the Nervous System* (Academic Press, New York, 1975). [20] W. Freeman and W. Schneider, *Psychophysiology* 19, 44 (1982).

[21] G. Viana di Prisco and W. Freeman, *Behav. Neurosci.* 99, 964 (1985).

[22] W. Freeman and G. Viana di Prisco, *Behav. Neurosci.* 100, 753 (1986).

[23] W. Freeman, *Biol. Cybern.* 56, 139 (1987).

[24] Y. Yao, W. Freeman and B. Burke, *Neural Networks* 3, 153 (1990).

[25] W. Freeman, in *The Ubiquity of Chaos,* ed. by S. Krasner (American Association the Advancement of Science, Washington, DC, 1990) p. 47.

[26] M. Merleau–Ponty, *Phenomenologie de la* Perception (Presses Universitaires de France, Paris, 1962).

[27] A. Garfinkel, *Behav. Brain Sci.* 10, 178 (1987).

[28] A. Garfinkel, *Forms of Explanation* (Yale University Press, New Haven, 1981).

[29] J. Crutchfield, J. Farmer, N. Packard and R. Shaw, Sci. *Am.* 256, 46 (1987).

[30] S. Rose, *Molecules and Minds* (Open University Press, Milton Keynes, 1987).

[31] M. Merleau–Ponty, *La Structure du Comportement* (Presses Universitaires de France, Paris, 1942).

[32] P. Churchland, *Neurophilosophy* (MIT Press, Cambridge, Massachusetts, 1986).