Mind/Brain Science: Neuroscience on philosophy of mind. *Chapter 7: Lepore E, van Gulick R (eds.) John Searle and his Critics. Oxford UK: Blackwell. pp. 115-127. Freeman WJ, Skarda CA (1990)*

Mind/Brain Science: Neuroscience on Philosophy of Mind

WALTER J. FREEMAN

CHRISTINE A. SKARDA

1. Introduction

With the advent of behaviorism it became fashionable to claim that neuroscience is irrelevant to cognitive science, and since the introduction of the computer, philosophers of a functionalist persuasion have pointed to the symbolic processing model to support this view. Fortunately, this line of thinking did not stop most of the minds/brains in neuroscience from pursuing research that ran directly counter to the functionalist doctrine, Neuroscientists never believed that knowing how the brain functions is irrelevant to cognitive theory: functionalism looked nice on paper to those who knew nothing about how brains work, but it was based on profoundly anti–biological assumptions which, far from encouraging interdisciplinary exchange, prompted most neuroscientists to ignore what was going on in cognitive science. Today, as functionalism gives way to "neurophilosophy" (Churchland, 1986) and "neural network" models replace some forms of symbol–based processing in artificial intelligence, cognitive and neuroscientists should look at what is being said by their colleagues in both disciplines.

John Searle's work in philosophy of mind has distinguished itself for many reasons, chief among these being that he never accepted functionalism and, unlike many philosophers, always enthusiastically embraced neuroscience. We suggest that his contributions deserve serious consideration, in turn, by neuroscientists. Of special interest in this respect are Searle's views on the mind/brain. Our comments focus on four of these: (1) the relation between brains and computers; (2) the issue of levels of description; (3) Searle's model for interpreting how the mind/brain functions; and (4) his characterization of perceptual states. Recent findings in neuroscience corroborate many of Searle's claims and support his view that there is no mind/body problem except in the minds of some philosophers. But our research also suggests that some traditional, time–worn and no longer valid physiological concepts persist in Searle's characterization of how the mind/brain works which lead to problems that neuroscience may help to solve. We address these issues in an attempt to contribute to a future mind/brain science that is being envisaged by many researchers today, including Searle.

2. Brains and Computers

Searle has presented a forceful and convincing argument against the view that intelligence is formal symbol manipulation (Searle, 1980). This view, which we also believe to be false, says that the brain produces intelligent behavior because at the functional level of description the brain is a device formally manipulating symbols according to rules. This view characterizes classical cognitivism; it is still not uncommon to read that a model that does without symbols and rules at the functional level is "noncognitive" (Earle, 1987).

Evidence from our laboratory indicates that brain functioning does not resemble the rule–driven symbol–manipulating processes characteristic of digital computers (Skarda and Freeman, 1987). Electroencephalogram (EEG) research on preattentive sensory processing taking place in the olfactory bulb suggests that brains use the functional architecture of distributed, self–organizing networks similar in many ways to some types of present–day connectionist models. The essential feature of connectionist or so–called "neural net" models is that a distributed system of interacting simple elements can produce intelligent behavior without the rules and programs that were previously thought to be required. Learning takes place by strengthening and weakening connection strengths between units in the network in parallel. When the network is activated each unit computes its own level of activity in terms of input from other units and a predetermined threshold value. The global pattern of activity resulting from these simultaneous, independent, parallel computations constitutes the state of the system at each moment.

To support Searle's claim with an appeal to connectionist systems, however, we believe it is important to distinguish two camps of connectionist models, one typified by so–called PDP models (Hinton, 1985;

Rumelhart et al., 1986), the other characterized by self–organizing dynamical systems (Amari, 1983; Anderson et al., 1977; Freeman, 1987; Grossberg, 1981; Hopfield, 1982; Kohonen, 1984). Not every connectionist network can be said to operate in the absence of symbols or symbol–like elements. Systems that fall within the PDP class of connectionist models use globally distributed dynamics, but there is a sense in which this class of systems still uses internal "representations" in the production of behavior. Systems like these, which rely on feed–forward connectivity and back propagation for error correction, have their "goals" externally imposed: they require a "teacher" or set of correct answers to be instilled by the system's operator. These answers are paradigmatic patterns with reference to which the output of the system is corrected via back propagation and error correction. The teacher used by such systems may not be contained in a program, but it nonetheless functions in the same way that an internal representation does in conventional computers. Self–organizing dynamic systems, because of dense local feedback connections, do not require or use teachers. No matching or comparison takes place such as by correlation or completion, and no archetypal set patterns are placed by an external operator into the system as its goals.

Our data indicate that in the olfactory bulb, learning consists in a self–organizing process and not in the manipulation of symbols according to rules: it is the selective strengthening of excitatory connections among neurons (mitral cells) in the bulb leading to the constitution of a nerve cell assembly (NCA) and ultimately to a change of state from local activity in the bulb to a self–organized form of activity globally distributed over the entire set of bulbar neurons that can be related with a particular odor and response. Memory for an odor consists in the set of strengthened excitatory connections of the NCA that, when activated with input, possesses the tendency to produce a spatial activity pattern characteristic of a given odor. These are not mechanisms used by conventional digital computers. No program–specified rule is imposed on olfactory input, the activity is self–organized, there is no central processor, and learning and memory are distributed throughout the system.

On the basis of our data we have suggested that brains do not work in the way in which everyone expected and that new conceptual apparatus will be required to characterize adequately the self–organized (see section 4 below) neural processing that we have observed. With the introduction of alternative models, the identification of symbol processing with intelligence appears to be an oversimplification at best, and we, along with Searle, have argued that it is misleading for both cognitive and neuroscientists to think in these terms when attempting to understand brain function (Freeman and Skarda, 1985; Skarda, 1986; Skarda and Freeman, 1987).

3. Levels of Description and the Neuron Doctrine

We agree with Searle that "[pains] and other mental phenomena just are features of the brain and perhaps the rest of the central nervous system," and that the important requirement for understanding this relationship is the distinction between micro– and macrolevels of neural functioning, The brain produces intelligent behavior, but the crucial question for neuroscientists has always been which level of functioning is relevant for explanations of this behavior? Our research has led us to break with a foundational concept of contemporary research on the nervous system, the "neuron doctrine," that we once accepted in company with the majority of our colleagues, but which we now see as mistaken and as a source of misunderstanding in attempts to comprehend the brain as the organ of behavior (Freeman, 1984).

The nineteenth–century witnessed two main revolutions in biology. The best known is the introduction of Darwinian doctrine; the other is the cellular doctrine of Rudolf Virchow (187 1), by which we came to understand that the functional basis of all life is the cell. The teaching and practice of physiology and medicine today are entirely based on this doctrine, and it is undoubtedly correct to say that the great majority of janitorial functions of neurons (processes of growth, repair, general chemical maintenance, metabolic fueling, disposal of wastes, and the modification of connections in molecular and membranal processes) take place and are best understood at the cellular level. The crucial question, however, is whether information used in the elaboration of goal–directed behavior and in controlling interactions with the environment is expressed in or operated on by neurons acting as individual summing and switching devices?

For nearly a century the common answer has been "Of course it is!" The principle reason for believing this has been the assumption that what counts in the production of behavior is the action potential, the most prominent signal generated by individual neurons (Barlow, 1972). Such a cell is called a "unit" with a "spike train," and it is commonly thought that the "message" or "information" transmitted by the cell consists in the temporal patterns of the spikes, in the manner that a message resides in the pattern of dots and dashes in a telegraph wire.

We note, however, that the great bulk of physiological studies cited to support this view come from anesthetized or paralyzed animals in circumstances in which goal–directed behavior is deliberately suppressed. When studies have been undertaken in normally behaving animals, the variability of the observed unit activity has been so great as to be uninterpretable, so that researchers have resorted to Averaging their data in such a way that the time base is locked to a stimulus or response. This procedure culls from the data an extremely small fraction of the variance locked to an external event and discards the overwhelming majority of the data as "noise."

Our studies of the behavior of single neurons in respect to the tens and hundreds of thousands of their neighbors indicates that researchers have been searching for neural "information" at the wrong level. While the activity of single cells appears to be largely unpredictable and noisy, the mass of cells cooperates to produce a coherent pattern that can be reliably related with a particular stimulus. In studies of the olfactory systems of small mammals the results are unequivocal: the information expressed by neurons that is related with the behavior of animals exists in the cooperative activity of many millions of neurons and not in the favored few. In the first stage of the system where odor discrimination takes place, the olfactory bulb, there is no evidence that the signaling by the bulb to the rest of the brain is by any small number of neurons unique for each odor. On the contrary, for every discriminable odor *every* neuron in the bulb participates. We conclude that it is incorrect to say that the behavioral information exists in the activity of single neurons; it cannot be observed there. This information exists in patterns of activity that are distributed concomitantly and continuously over tens and hundred of millions of neurons.

The main reason that the cellular doctrine holds so well for most cells but fails for neurons is that neurons are uniquely different from all other types of cells: neurons are involved constantly in widespread activity with other neurons, some at great distances. Each neuron transmits to thousands of others in its vicinity, and in some cases for astonishingly great distances, often thousands of times greater than the diameter of its own axon. It also receives from many neurons in its surround, on the order of 1,000 to 10,000. The key to its interaction lies in this: by virtue of its membrane each neuron has a relatively high degree of automomy, yet all that it does is felt by many other neurons, and those neurons provide the environment to which it responds. These widespread actions and reactions lead to the emergence of the constant, ceaseless, ever–fluctuating activity of masses of nerve cells that start talking and never stop. In fact, if they do stop, because they are cut off from interaction with their neighbors, they soon atrophy and die.

Interconnections among neurons are of two main types: local and long range, corresponding to Golgi II and Golgi I neurons. Local connections within areas of the cortex and subcortical nuclei form the basis for interactive masses that create and maintain fluctuating patterns of activity over the entire spatial extent of each mass. Distant connections serve to transmit patterns from one local mass to another and back again, almost always in reciprocal pairs again subserving interaction and not merely action. These reciprocal connections form. the basis for feedback loops that are the hallmark of brain structure and function. The overwhelming majority of past and contemporary studies of the brain have omitted these feedback pathways, expressing brain dynamics in terms of flow diagrams that show inferred connections as all being feedforward. But with the introduction of new mathematics and technology it has become possible to simulate the dynamics of neural masses in the brain with large sets of differential equations and to display the solutions as graphic patterns.

These findings have encouraged us to break with the "reductionist" view that the behavior of a system can be explained in terms of the properties and relationships between individual components that constitute the system. Another feature of brain function is also important in this connection: brains use chaotic dynamics in the production of behavior (Skarda and Freeman, 1987). There is a simple recipe for creating this kind of dynamic activity, and it is one that brains use. (1) Assemble together a large number of distinct elements such as molecules, neurons, or even people and allow each to transmit to and receive from many others in the group information, matter, and energy. (2) Specify that the relationship between input and output for each element be nonlinear. This means that if the input is increased in small steps from a low level to a high level, the output changes but not in direct: proportion to the input. For example, in the brain if a neuron is stimulated weakly it may release currents but not action potentials, but once a certain threshold of input is reached it will fire. This sudden "jump" or change of behavioral state reflects nonlinearity, i.e., a disproportionately large increase in output for a given input. (3) Make sure that the system is "open", i.e., there is a ready supply of energy, food, blood, and a good disposal system for heat and other wastes. (4) Finally, turn up the temperature, e.g., apply heat uniformly, or infuse an excitatory chemical. Given these conditions, something interesting is likely to happen, and it is likely to be unpredictable.

With some initial conditions and arrangements there will emerge new and sometimes fascinating patterns of behavior. Some may congeal and terminate the experiment; others may move, rotate, and reform periodically; but others may dissolve into ceaseless activity without discernable spatial and temporal structure that resembles "noise." This ceaseless activity is chaos.

"Chaos," in its traditional meaning, refers to complete disorder, the formless void from which all order in the universe arose. The new mathematical meaning for the term refers to activity that appears random, but is not. It is deterministic, in the sense that it can be reliably simulated by solving sets of coupled nonlinear ordinary differential equations or generated by building a system to certain specifications and putting energy into it. Chaos exists in many forms and degrees (see e.g., Crutchfield et al., 1987). It can be reproduced with high precision if the initial conditions are identical on repeated runs, but it is unpredictable if new initial conditions are used. Chaos has precisely definable qualities and has relatively small degrees of freedom, meaning that its dynamics can be described by relatively few variables and therefore dimensions or coordinate axes. Most important, it can be turned on and off virtually instantaneously as with a switch.

The physiological basis for the view that brains employ chaotic dynamics involves a hypothesis on the way in which synaptic strengths change during learning under reinforcement. Essentially, when two neurons fire together, i.e., when the action potential of the presynaptic neuron excites the postsynaptic. neuron and generates an action potential, the synapse that connects them is strengthened. This is a general form of the so–called Hebb rule of synaptic learning (Hebb, 1949; Viana Di Prisco, 1984). When many interconnected neurons within a mass of neurons fire together in pairs over repeated stimuli, the selectively co–active neurons are joined together into a network of strengthened connections. This conjoined set is called a nerve cell assembly (NCA), and we believe it is the basis for perception and learning in the nervous system. When perception takes place it is expressed in a reproducible and identifiable pattern of activity that is mediated by the NCA.

When a novel stimulus is presented under reinforcement, a new NCA forms during the first few presentations. However, the background activity state of neural activity that goes on in the absence of the new reinforced stimulus should not be patterned, because this would drive the system into an already existing pattern of activity associated with another stimulus. The kind of activity that is required in order that a new pattern can emerge must be unpatterned yet controllable. This is chaos. It is generated by the nervous system in the presence of a novel stimulus so that the neurons have activity by which the Hebb rule can operate. Thereby a novel pattern can emerge from the chaotic state, and the system is not forced into preexisting "grooves" or patterns of activity. Chaos has been identified in many areas of the brain (Babloyantz and Destexhe, 1986; Nicolis and Tsuda, 1985; Freeman and Viana Di Prisco, 1986; Garfinkel, 1983; Skarda and Freeman, 1987) with the implication that it may provide the basis for flexibility, adaptiveness, and the trial–and–error coping that make possible the nervous system's interaction with an unpredictable and ever–changing environment.

The observation that brains employ chaos to produce behavior is important in the present discussion because phenomena that are chaotic preclude long–term predictions. Chaotic behavior emerges from the nonlinear interaction of its parts, and global behavior in the system cannot be reduced to or deduced from knowledge about the characteristics and interactions among individual components (Crutchfield et al., 1987). As we have indicated, it is not individual neurons and their activities that explain or cause behavior; it is rather the activity produced by masses of neurons that self–organize to produce new global forms of behavior. As Searle indicates, "just as we need the micro–macro distinction for any physical system, so for the same reasons we need the micro–macro distinction for the brain" (Searle, 1984).

But our agreement with Searle is predicated on an important caveat, namely, that he means the same thing we do by micro/macro descriptions, It isn't always clear that he does mean what we mean. We agree with Searle when he says that while it makes sense to say that a particular organism with its nervous system is experiencing a given stimulus, it does not make sense to say for any particular neuron in that brain that it is experiencing that stimulus (Searle, 1984). The brain, on our view, gives rise to emergent neural phenomena that are responsible for and explain behavior and that are not reducible to the features and relations of its component parts. But sometimes Searle appears to make a further claim, another kind of level distinction, that is not encompassed by the one we have discussed. This distinction is between a microlevel that includes all levels of neuronal processing and a macrolevel of purely mental processes. Searle says, "At the higher level of description, the intention to raise my arm causes the movement of the arm. But at the lower level of description, a series of neuron firings starts a chain of events that results in the contraction of the muscles" (Searle 1984).

There are two things to say about this claim. First, it by–passes the role played by global neural activity and conflates the global neural and mental levels. Second, although Searle argues against dualism, he seems to argue here for a level of activity that plays a causal role but is not physiological, a specifically "mental" level. This is philosophically appealing but lacks biological sense. The reason is that a,, physiologists we cannot make strict causal inferences from the level of neurons to that of neural mass actions (see above); a fortiori, we cannot impute cause and effect between the global neural and mental levels. Quite apart from the classical problems concerning causality raised by Hume, Kant, Whitehead, and others, our modern conceptions of feedback necessarily introduce ambiguity and indeterminacy. These are endlessly compounded in our efforts to comprehend distributed networks with large numbers of feedback loops. Already we can build models of brain parts that function in some of the ways that brains do. We can describe then and largely (not entirely) control them, but we cannot explain how they work (Hopfield and Tank, 1986). It appears to us likely that during the next decade or so, machines will be constructed that will display useful traits heretofore restricted to biologic intelligence, and the irony will be that we will be unable to understand their processes in causal terms. The problem this raises for Searle is that where he want a causal explanation there isn't one to be had.

4. The Reflex Model of Behavior

Our data have forced us to reject yet another foundational concept implicit in much of contemporary physiological research: the reflex doctrine. Physiologists an psychologists who work with animals often have the illusion that they control their behavior by use of reinforcement. This belief is based on the model of physiological functioning developed to explain reflex behaviors (Sherrington, 1906) and on thee feedforward models that experimentalists use to explain how it is that a conditioned stimulus (CS) that is paired with an unconditioned stimulus (UCS) will elicit conditioned response (CR) preceding the unconditioned response (UCR). The presumption is that all behavior can be expressed as a sum of responses to stimuli, a view that includes and is ultimately derived from such fundamental behaviors as the slaking of thirst, the satisfaction of hunger, and the titillation of sex.

What experimentalists have failed to note is the essential fact that in the typical experiment it is the animal that is controlling their behaviors: researchers spend, as they should have spent, small fortunes on the care, feeding, and housing of the subjects; they tailor the equipment and tasks to the capabilities of the species, familiarize and train them, and then sit waiting for them to deign to stop eating licking, grooming, or just looking around long enough for the experimenter to get is a CS for a controlled trial; all this can go on for weeks. What is lost in all this is the fact that these animals are continually producing behaviors from within by anticipating external stimuli to guide or pace their actions. These behaviors express internally generated activity of the nervous system and are not determinist responses to stimuli.

There is a term to describe such internally generated activity. We say that this type of neural activity is "self–organizing". We see self–organizing activity in many inanimate systems around us: in the emergence of patterns of clouds in a previous clear blue sky, in the bubbles that form at the bottom of a heated pan of water, and in the formation of drops of water from a leaky faucet. In the biological sphere we see self–organized dynamics in the earliest stages of development of structure and function from the fertilized ovum on into the growing embryo. Recent findings indicate that behavior can arise from within the system by self–organized patterns of neural activity in interconnected masses of the brain, and that it is not simply the sum of conditioned or unconditioned responses to stimuli (Freeman and Skarda, 1985; Skarda and Freeman, 1987).

In the olfactory bulb, self–organized patterned activity is essential for odor recognition and discrimination, processes that cannot be explained in terms of the reflex model. Molecules of the odorant fall onto the cilia of olfactory receptor cells and, following capture, excite a small subset of receptors that are selectively sensitive to the odorant. This step is called transduction. In the next step of forward excitation the receptors send action potentials into the olfactory bulb that excite the large projection neurons known as mitral cells. Those that have been previously excited by the odorant under reinforcement, and have had their synaptic interconnections strengthened by the learning process described earlier, preferentially excite each other. This subset of mitral cells forms a nerve cell assembly of strongly interconnected cells. In the third stage of feedback interaction the initial excitation spreads like wildfire through the assembly. Those other mitral cells that have been excited, although not under reinforcement, but that receive excitation from background odorants, excite the bulb further. Both kinds of excitation serve to increase the sensitivity of bulbar neurons to each other and therefore to increase intrabulbar crosstalk. When a critical threshold is reached the fourth step occurs. This is a state change via a self–organized dynamic process in which the entire bulb goes from chaotic activity to an internally generated burst of oscillatory activity that we can correlate with a particular odorant. This globally distributed, self–organized, stereotypical pattern of activity in the bulb is the one that we hypothesize constitutes odor recognition and memory for the rest of the system; this is the pattern that is made available to the rest of the brain and that is behaviorable relevant for the kind of relations that we usually associate with learning and memory (Skarda and Freeman, 1988).

This feature of nervous system functioning, i.e., that it creates something that did not exist before its interaction with the environment, has not gone unnoticed in the past, but before the introduction of the mathematical theory of nonlinear dynamics and computers there seemed no viable alternative explanatory framework to stimulus –response determinism in experimental brain science. Early on, Sherrington, who developed the synaptic basis for the concept of the reflex are, uncovered an internally generated neural process that acted on stimulus input but was not caused by it (Sherrington, 193 1). Sherrington suggested that this "central excitatory (inhibitory) state" was internally generated by a mass of neurons that then serve actively to coordinate and integrate reflexes that, in themselves, were produced passively by a stimulus.

But physiologists were not alone in recognizing that neural dynamics are more than reactions to external stimuli. Philosophers, too, have realized that the reflex model of physiological functioning is inadequate for explaining behavior. Dewey (1896) can be credited with probably the earliest attack on the notion. As Dewey put it, the stimulus "must already have one foot over the threshold, if it is ever to gain admittance" (Dewey, 1896), and this is provided by a internally motivated activity.

Later Merleau–Ponty (1942) argued that behavior, and hence a neural state generated from within the organism, is the first cause of all stimulation citing the very same data used to support reflex theory by Pavlov (1927) and Sherrington. Perception, on this view, begins within and is not imposed from without by the stimulus, If the opposite view based on the reflex theory dominates our thinking, this is not because we have been convinced by the data but because we interpret all the data to fit the reflex theory we have about how perception must work. At best, Merleau–Ponty suggested, the activity described by classical reflex theories is pathological.

> The reflex as it is defined in the classical conception does not represent the normal activity of the animal, but the reaction obtained from an organism when it is subjected to working as it were by means of detached parts ... [it is] characteristic not of the fundamental activity of the living being but of the experimental apparatus which we use for studying it. (Merfeau–Ponty, 1942)

Our data support Dewey's and Merleau–Ponty's positions. Only under reinforcement, i.e., only in animals who have been motivated to expect certain odorants, does receptor activity initiated by the stimulus lead to changes in the patterned activity of the bulb, the generation of a NCA, etc. In the neural system, if previous, internally generated activity has not laid the groundwork for interaction with the environment, receptor stimulation cannot and will not lead to perception and/or behavior (see also section 5).

We mention this because very often it is overlooked by contemporary philosophers interested in understanding the physiological basis of behavior (Skarda, 1986). The tendency is to adopt a reflex–based model and to generalize it for all behavior. We believe that this move is not only bad physiology; it is misleading for philosophical attempts to understand the biological basis of behavior and the mind/ brain relationship (see section 5 below and Skarda, 1986). We think that Searle falls into this trap by focusing on the neurophysiological basis of behaviors that are reflex in nature, e.g., the slaking of thirst or the contraction of muscles, and then adopting this model for all brain function. If reflexes were all that were needed to produce behavior as we know it, we would not find brains generating self–organized behavior and chaos; and certainly, by focusing attention on reflex behavior Searle misses, along with many others, the most distinctive feature of brain functioning in the production of goal–directed behavior, i.e., its self– organizing, creative dynamics. We encourage philosophers and cognitive theorists to take another look at what neurophysiologists have discovered about the neuronal processes responsible for behavior, and we hope that in the future more attention will be given to the excellent work done by distinguished thinkers like Dewey and Merleau–Ponty who pointed out long ago that there is more to behaving than reacting.

5. What's in Perceptual States?

In a recent book on the philosophy of mind Searle discusses what he terms the "intentional contents" of perceptual states (Searle, 1983). He describes these contents as follows:

vis exp (that x with certain features is before me and that x is causing this perceptual experience)

Whatever else is packed into this description, it implies that what counts as far as perception (in this case visual perception) is concerned, i.e., the essential features of the perceptual process, are an internal representation of the object perceived along with its features and the causal impact that the object has on the system perceiving it. What does neurophysiology have to say about Searle's description?

The first thing to note about Searle's description of perceptual states is that it incorporates the reflex model discussed above in section 4. Studies of physiological functioning, however, reveal that perceptual processing involves more interesting processes than those described by traditional reflex–based theory and feedforward processing have imagined. Evidence indicates that when an organism is trained to respond to a particular odorant a self–organized process in the bulb produces a spatially coherent activity state that can be modeled as a limit–cycle attractor, Such attractors mathematically represent the qualitative form of behavior exhibited by a system, in this case periodic or aperiodic (chaotic) behavior. Topologists say that the behavior is governed by a stable attractor if the system returns to the same form of behavior after it is perturbed and is allowed to settle. With each inhalation, after learning and in the presence of this odorant, this more ordered state repeatedly emerges from the chaotic background state, only to collapse back again with exhalation. A separate spatial pattern of periodic behavior forms for each learned odor given under, reinforcement. Each has a latent coexistence in that only one at a time can find expression or be realized and then only, as far as we have yet been able to detect, on presentation of its odorant under reinforcement (we have not sought to explore hallucinations). When the reinforcement contingency is changed in respect to any one odorant, or if a new odorant is added to the repertoire under reinforcement, all the spatial patterns undergo small changes during the process of learning. These changes do not occur in the olfactory bulb if there is no reinforcement or if the newly learned CS is not olfactory but visual or auditory. These spatial patterns are manifested in the bulbar EEG and their information content is measured by assaying the capacity of our measurements on these patterns to classify correctly the bursts in respect to what CS was given and what CR occurred on successive conditioning trials (Freeman, 1987).

Several features of the neural dynamics underlying perception are important to note. First, only when the odorant is reinforced leading to formation of a CR, i.e., only when the animal is motivated and the stimulus input has some meaning for the organism such that it acts on the stimulus, do odor–specific activity patterns form in the olfactory bulb. Presentation of odorants to the receptor surface in unmotivated subjects does not lead to any observable changes in the system. Second, the odor specific activity patterns are dependent on the behavioral response: changing the reinforcement contingency changes the activity patterns previously recorded. Finally, the internally generated odor–specific activity is context– dependent: introducing new reinforced odorant to the animal's repertoire leads to changes in the activity patterns of all previously learned odorants; in other words, adding a new odor under reinforcement introduces not only a quantitative change in the number of learned activity patterns, it also qualitatively alters each of the patterns previously learned.

How do these findings bear on Searle's claims regarding perception? As mentioned above, Searle's characterization of perceptual contents implies that what is important for the system are the object with its features and the object's causal impact on the perceiving subject. He states that '[t]he story *begins* with the assault of the photons on the photoreceptor cells of the retina, the familiar rods and cones" (Searle, 1983). Our data indicate that perception begins with an internally generated neural process that, by re–afference, lays the ground for processing of future receptor input. The neural activity patterns that we find related to perception are indicative of internal states that reflect reliable forms of interaction in a context. They are sensitive not simply to the presence of an odorant, or to the response, but to both in interaction and to the context of reinforced odorants in which this behavior is embedded. For Searle, perception is something that happens to the system when it is *acted upon* by an object; for us, perception is a process that occurs only when the organism *initiates interaction with* its environment.

In review we suggest that two things are missing in Searle's description: neural interaction and internal context. With respect to the first, causal impact on the system should not be confused with interaction in our sense.

As we have indicated, causal stimulation at the receptor level alone does not lead to the formation of odor–specific activity patterns in the bulb that are the basis of odor recognition and discrimination. Only when the odorant is reinforced, when it acquires behavioral significance in terms of a CR, does the internally generated pattern form. Thus, the story of perception cannot be told simply in terms of feedforward causation in which the object initiates neural changes leading to an internal perceptual state. What is missing here is recognition of the role played in perception by self–organizing neural processes and by the dense feedback among subsystems in the brain that allow the organism to initiate interaction with its environment.

Second, internal context is also part of the meaning of the activity pattern for the system: when a new odorant is added to the animal's repertoire *all* previously learned patterns undergo a change. Searle is aware that context plays a role in perception as in all intentional states. He terms this context the "Network" (Searle, 1983), and he claims that each state is embedded in this Network of other states, The problem is that nowhere *within* his description of the intentional content of perceptual states does he include a reference to this Network. For Searle, each perceptual state is located in a Network of other states, but the presence or absence of states in this Network does not effectively alter the content of any one. We have found, however, that neural dynamics are sensitive to, and hence the characterization of their content must include information on, interrelationships within a perceptual system or subsystem, because adding to this Network of perceptual states in the olfactory system leads to a change in the patterned activity of each and every odor–specific state. Thus, we support Searle's insight that perceptual states are always located in a network of states, but our evidence suggests that this network is *internal* to each state. We believe that Searle is on the right track, as usual, but the predominantly reflex–based conception of brain functioning that he has inherited from classical neuropsychology has stymied him. We suggest that brain theories couched in self–organizing nonlinear dynamics will provide the keys he and others need to solve the problem of explaining how intentionality can function with the physiochemical organism.

6. Conclusions

We are greatly encouraged by the attempts of contemporary philosophers of mind to include neuroscience in debates about the mind/body problem and the explanation of behavior. It is high time for a biologically grounded cooperative effort to understand the mind/brain, and Searle has done much to contribute to this project. We are eager to participate, for we believe that there will never be a coherent theory of the mind/brain and its role in behavior as long as philosophers ignore the brain and as long as neuroscientists shy away from the mind. We hope that our remarks on Searle's work illustrate the fruitfulness of cooperative efforts, both with constructive criticism and support based on our data. Without theories neuroscience has no way to explain its data, but without facts about how brains actually function philosophical theories tend to hang in the air. It is only taken together that theories and facts lead to insights and direct future research. We hope this is the way of the future.

REFERENCES

Amari, S. (1983) "Field theory of self–organizing neural nets," *IEEE Transactions on Systems, Man, and Cybernetics,* SMC–13, 5, pp. 741–8.

Anderson, J., Silverstein, J., Ritz, S. and Jones, R. (1977) "Distinctive features, categorical perception, and probability learning: Some applications of a neural model," *Psychological Review* 84, pp. 413–51.,

Babloyantz, A. and Destexhe, A. (1986) "Low–dimensional chaos in an instance of epilepsy," *Proceedings of the National Academy of Sciences of the United States of America* 83, pp. 3513–17.

Barlow, H. (1972) "Single units and sensation: A neuron doctrine for perceptual psychology?" *Perception 1,* pp. 371–94.

Churchland, P. (1986) *Neurophilosophy: Toward a Unified Science of the Mind–Brain* (MIT Press/ Bradford: Cambridge, Mass.).

Crutchfield, J., Farmer, J. Packard, N. and Shaw, R. (1987) "Chaos," *Scientific American* 256, pp. 46–57.

Dewey, J. (1896) "The reflex arc concept in psychology," *Psychological Review* 3, pp. 357–70.

Earle, D. (1987) "On the differences between cognitive and noncognitive systems," *Behavioral and Brain Sciences 10, pp.* 177–8, Freeman, W. (1984) "Premises in neurophysiological studies of learning," in *Neurobiology of Learning and Memory,* ed. G. Lynch, J. MeGaugh, and N. Weinberger (Guilford Press: New York, NY).

Freeman, W. (1987) "Analytic techniques used in the search for the physiological basis of the EEG," in *Handbook of Electroencephalography and Clinical Neurophysiology,* ed. A. Gevins and A. Remond (Elsevier: New York, NY), vol. 3A, part 2, chapter 18.

Freeman, W. and Skarda, C. (1985) "Spatial EEG patterns, nonlinear dynamics and perception: The neo–Sherringtonian view," *Brain Research Reviews 10, pp.* 147–75.

Freeman, W. and Viana Di Prisco, G. (1986) "EEG spatial pattern differences with discriminated odors manifest chaotic and limit cycle attractors in olfactory bulb of rabbits," in *Brain Theory,* ed. G. Palm (Springer Verlag Berlin).

Garfinkel, A. (1983) "A mathematics for physiology," *American Journal of Physiology* 245: *Regulatory, Integrative and Comparative Physiology* 14: R455–66.

Grossberg, S. (1981) "Adaptive resonance in development, perception, and cognition," in *Mathematical Psychology and Psychophysiology,* ed. S. Grossberg (American Mathematical Society: Providence, R.I.).

Hebb, D. (1949) *The Organization of Behavior* (Wiley: New York, NY).

Hinton, G. (1985) "Learning in parallel networks," *Byte 10, p.* 265.

Hopfield, J. (1982) "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences of the United States* of *America* 79, pp. 2554–8.

Hopfield, J. and Tank, D. (1986) "Computing with neural circuits: A model," *Science* 233, pp. 625–33.

Kohonen, T. (1984) *Self– Organization and Associative Memory* (Springer Verlag: Berlin).

Merleau–Ponty, M. (1942) *La Structure du Comportement* (Presses Universitaries de France: Paris).

Nicolis, J. and Tsuda, 1. (1985) "Chaotic dynamics of information processing. The 'magic number seven plus–minus two' revisited," *Bulletin of Mathematical Biology* 47, pp. 343–65.

Pavlov, 1. (1927) *Conditioned Reflexes* (Oxford University Press: Oxford).

Rumelhart, D., McClelland, J. and PDP Research Group (1986) *Parallel Distributed Processing: Explorations in the Microstructures* of *Cognition. Vol. L Foundations.* (MIT Press/Bradford: Cambridge, Mass.).

Searle, J. (1980) "Minds, brains, and programs," *Behavioral and Brain Sciences 3, pp.* 417–57.

Searle, J. (1983) *Intentionality* (Cambridge University Press: Cambridge).

Searle, J. (1984) *Minds, Brains, and Science* (Harvard University Press: Cambridge, Mass.).

Sherrington, C. (1906) *The Integrative Action of the Nervous System* (Yale University Press: New Haven, Conn.).

Sherrington, C. (1931) "Quantitative management of contraction in lower level coordination," Hughings Jackson Lecture of 193 1, in *The Selected Writings of Sir Charles Shemington,* ed. D. Denny–Brown (P. B. Hoeber: New York, NY, 1940).

Skarda, C. (1986) "Explaining behavior: Bringing the brain back in," *Inquiry* 29, pp, 187–202.

Skarda, C. and Freeman, W. (1987) "How brains make chaos in order to make sense of the world," *Behavioral and Brain Sciences 10, pp.* 161 – 73.

Skarda, C. and Freeman, W. (1988) "Research of neural dynamics: Implications for models of learning and memory," in *Systems with Learning and Memory Abilities,* ed. J. Delacour and J. C. Levy (Elsevier: New York, NY).

Viana Di Prisco, G. (1984) "Hebb synaptic plasticity," *Progress in Neurobiology* 22, pp. 89–102.

Virchow, R. (1871) *Die Cellula?pathologie in ihrer Begriindung auf Physiologie und pathologische Gewebelehre (A*. Hirschwald: Berlin).

END