

How brains make chaos in order to make sense of the world

Christine A. Skarda

CREA, Ecole Polytechnique, 75005 Paris, France

Walter J. Freeman

Department of Physiology-Anatomy, University of California, Berkeley, Calif. 94720

Abstract: Recent “connectionist” models provide a new explanatory alternative to the digital computer as a model for brain function. Evidence from our EEG research on the olfactory bulb suggests that the brain may indeed use computational mechanisms like those found in connectionist models. In the present paper we discuss our data and develop a model to describe the neural dynamics responsible for odor recognition and discrimination. The results indicate the existence of sensory- and motor-specific information in the spatial dimension of EEG activity and call for new physiological metaphors and techniques of analysis. Special emphasis is placed in our model on chaotic neural activity. We hypothesize that chaotic behavior serves as the essential ground state for the neural perceptual apparatus, and we propose a mechanism for acquiring new forms of patterned activity corresponding to new learned odors. Finally, some of the implications of our neural model for behavioral theories are briefly discussed. Our research, in concert with the connectionist work, encourages a reevaluation of explanatory models that are based only on the digital computer metaphor.

Keywords: brain theory; chaos; cognitivism; connectionism; EEG; nonlinear dynamics; olfaction; perception; sensation

1. Introduction

To understand brain function we need to know how the sensory systems process their information. Recent connectionist models provide an interesting explanatory alternative to earlier information-processing models based on the digital computer that viewed neurons as two-state logical decision elements organized into networks to compute simple Boolean functions. In the present article we outline the results of experiments in our laboratory that demonstrate the existence of sensory- and motor-specific information in the spatial dimension of EEG activity in the central nervous system. On the basis of our data we develop an explanatory model of the neural states responsible for sensory encoding; this model departs significantly from alternatives patterned after digital computers and it converges with recent connectionist models in the computational principles it uses. We suggest, however, that brains rely on mechanisms not found in other models; we propose four such mechanisms that may be necessary to solve problems critical to the efficient functioning and survival of any system that has to behave adaptively in an environment subject to unpredictable and often violent fluctuations.

Special emphasis is placed in our model on “chaotic” brain activity.¹ We propose that the brain relies on chaotic as opposed to steady or random activity for several purposes: Chaos constitutes the basic form of collective neural activity for all perceptual processes and functions as a controlled source of noise, as a means to ensure

continual access to previously learned sensory patterns, and as the means for learning new sensory patterns.

2. Methodological considerations

How does a sensory system process information? Models based on the digital computer define computation as a physical operation governed by the substates of the parts of the system as defined by rules operating on symbol tokens in virtue of their formal syntactic structure corresponding to real physical differences in the system. The formal elements or symbols are required to be discrete – that is, context independent; each distinct semantic property must be associated with a distinct physical property (Pylyshyn 1984, pp. 50, 74).

For many years physiologists have applied the computational model when interpreting their data. Thus, they found that the “code” of peripheral sensory systems is based on “labeled lines” (Bullock & Horridge 1965, p. 274); the quality of a stimulus is conveyed by the selection of one or more axons from the immense number available, and the intensity is conveyed by the number of action potentials per unit time on each axon. This model worked for peripheral motor systems and for some parts of central nervous systems, to the extent that “feature detector” and “command” neurons could be identified. However, the search for this kind of information-processing scheme in the case of central associative functions has not been successful (Barlow 1972; Perkel & Bullock 1968).

Our attempt to understand information processing in olfaction was based on three premises. (1) When an animal that is conditioned to discriminate between two odorant stimuli inhales one of them (a conditioned stimulus [CS]) and then responds correctly (with a conditioned response [CR]), there will exist, somewhere and for some time during the interval between the onsets of the CS and CR, some odor-specific information in the olfactory bulb to serve as the basis for the correct CR. (2) This information will be encoded in the form of a space-time pattern of neural activity for each odorant CR. (3) These patterns will be manifested, however indirectly, in the electroencephalographic (EEG) potentials recorded from the bulbar surface. After 12 years of search we at last identified some of the postulated patterns (Freeman & Viana Di Prisco 1986a). The results were beyond surprising; they took us so far outside the range of our previous expectations that we had no physiological metaphors with which to pin them down, and we had to draw on some new and fascinating fields of mathematics and physics in order to understand their implications.

In principle the experiments were simple. Thirsty rabbits were conditioned (Viana Di Prisco & Freeman 1985) to lick (CR+) in response to an odorant (CS+) followed after 2 seconds by delivery of water, and merely to sniff (CR-) in response to an unreinforced odorant (CS-). Each rabbit had an array of 64 electrodes implanted permanently onto the lateral surface of the left olfactory bulb. The 64 EEG traces were amplified, filtered, and measured in brief time epochs within each trial; when made with adequate safeguards (Freeman 1987b) these measurements from collections of trials served to classify EEG epochs into groups both with respect to CSs and with respect to CRs. The odorant-specific information was found to exist in the spatial patterns of the amplitude of the waveform of an oscillation of EEG potential that was common to all 64 channels and, by inference, to the entire bulb. We concluded that every neuron in the bulb participated in every olfactory discriminative response because they all participated in the oscillation. All that distinguished one odorant EEG pattern from another was the spatial configuration of the average intensity over an event time window at the common frequency, in the manner that patterns of monochromatic light are distinguished from each other by shades of gray. Local variations in phase, amplitude modulation, frequency modulation, and other aspects of the 64 traces were not found to contain odorant-specific information.

With regard to our first premise (that odor-specific information must exist in the bulb), we chose to study the olfactory system because it is the simplest and phylogenetically the most stable and representative of sensory systems, is the best understood in its structure and function, and can be studied in its earlier stages without directly involving the brain stem and thalamus. We selected the rabbit because its head is sufficiently large to support the electrical connectors needed for chronic implantation and recording from 64 channels, yet the bulb is sufficiently small so that the electrode array forms a window covering a substantial portion of its surface area (20% in the rabbit, as opposed to 6% in the cat; Freeman 1978). We used appetitive conditioning so as to have distinguishable behavioral responses from each ani-

mal: licking with or without sniffing to the CS+ (CR+) and sniffing alone to the CS- (CR-). We found that high relative frequencies of occurrence for these auto-shaped CRs (naturally occurring motor activity patterns) emerged within a very few trials in the first session, that they were stably maintained for numerous sessions, and that they were subject to quantitative assay with ease and reliability (Freeman 1981; Viana Di Prisco & Freeman 1985).

3. Neurophysiological results

3.1. Spatial analysis of neural activity. With regard to our second premise (that the odor-specific information is encoded as space-time patterns of activity), the set of chemoreceptor neurons in the nose and the set of mitral cells in the bulb to which they send their axons (Figure 1) both exist in the form of a sheet. Evidence from measurements of receptor unit activity, the electro-olfactogram, and odorant absorption to the mucosa upon stimulation with odorants show that receptor cells sensitive to a particular odorant are clustered nonuniformly in density in the mucosa, and that their spatial patterns of activation differ for differing odorants (reviewed in Moulton 1976:

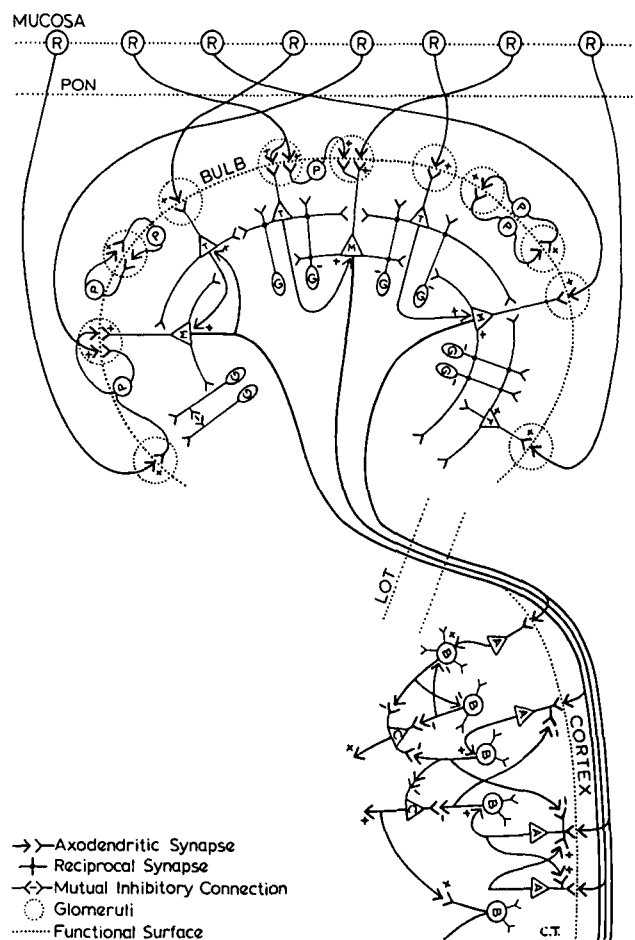


Figure 1. Schematic diagram of the main cell types and their interconnections in the olfactory bulb and prepyriform cortex: R, receptor; PON, primary olfactory nerve; LOT, lateral olfactory tract; M, mitral cell; G, granule cell; P, periglomerular cell; A, superficial pyramidal cell; B, granule cell; C, deep pyramidal cell. From Freeman (1972).

Freeman & Skarda 1985). The projection of the primary olfactory nerve (PON) onto the bulb has a degree of topographic order. Studies with 2-deoxyglucose (2-DOG) accumulation in the bulb after 45 minutes of exposure to an odorant show uneven clustering of dense patches in the outer (glomerular) layer of the bulb, indicating that a spatial pattern of receptor activity may result in a spatial pattern of neural activity in the bulb, which might in turn transmit odor-specific information to the olfactory cortex. However, metabolic studies cannot reveal the dynamic form of that neural activity in time periods on the order of 0.1 sec.

With regard to our third premise (that the odor-specific information is manifested in the EEG), single neurons in the receptor layer, bulb, and cortex respond selectively to test arrays of odorants at various concentrations. The variability and overlap of response profiles involving multiple odorants are high at all steps of the olfactory system, there being no indication that more centrally located neurons are more "narrowly tuned" to odorants than are receptors. The number and even the existence of "primary odors" analogous to colors or tastes are unknown.

Our early attempts to demonstrate spatial patterns of bulbar unit activity in responses to odorants were based on simultaneous multi-unit extracellular recording from 10 microelectrodes; the spatial sample was too small, and the time required to collect a sample (several minutes) was much too long. We turned to EEG recording from the bulbar surface because we had found a close statistical relationship in time and space between the amplitude of the EEG potential at selected points on the bulbar surface and the firing rates of mitral and tufted cells located at depths of several hundred microns below those points. That is, the surface EEG (Figure 2), consisting largely of extracellular compound postsynaptic potentials of granule cells, the dominant inhibitory interneurons deep within the bulb (Figure 1), provided indirect access to a spatial image of the locally averaged mitral cell activity patterns that constituted the bulbar output to the olfactory cortex. The theory and experimental evidence for this inference, including volume conductor theory and studies of the dynamics of bulbar neurons, have been compiled in a monograph (Freeman 1975) to which the interested reader is referred.

Measurements of the spatial spectrum of the bulbar EEG (Freeman 1980; Freeman & Baird, in press) were used to fix the optimal intervals between electrodes in arrays (the spatial digitizing increment) at 0.5 mm, corresponding to a Nyquist frequency of 1.0 c/mm. An 8×8 array gave a "window" onto the bulb of about 3.5×3.5 mm, given the restriction to 64 channels. Measurements of the temporal spectrum of the rabbit EEG indicated that the range of greatest interest was 20–90 Hz. Filters were set at 10 and 160 Hz; the temporal digitizing increment at 2 msec gave a Nyquist frequency of 250 Hz. A fixed duration of 76 msec was adopted as the minimum for the bulbar response on single inhalations, so that measurement of a single unaveraged event upon inhalation of an odorant or the background air for control consisted of 64×38 time values, digitized at 12 bits with retention of the 8 most significant bits. Each trial yielded 3 control events and 3 test odor events. Each session yielded 10 CS+ and 10 CS- trials, constituting 120




Figure 2. Four classes of states are identified for the olfactory system from EEG traces. Fluctuations are suppressed under deep anesthesia (lowest trace). In waking but unmotivated animals the amplitude is low and the trace is irregular and unpredictable. Under motivation the irregular activity is interrupted by brief oscillatory bursts following activation of the olfactory bulb by receptors on inhalation. Under several seconds of intense electrical stimulation of the LOT (top trace) an epileptic seizure is released. It is initiated after the failure of excitatory input transmission as shown by the decreasing responses at left to the last 5 pulses of the stimulus train. The seizure spike train then progressively emerges from a relatively quiet post-stimulus state. From Freeman (1987a).

events. The data base for the study comprised 18 sessions with each of 5 rabbits after a familiarization period.

Acquisition of these data required 64 preamplifiers, a high-speed multiplexer and ADC, and a dedicated computer (Perkin Elmer 3220) and disc. The limiting factor on data acquisition proved to be the core-to-disc data transfer rate with double buffering during the 6-second trial periods. Procedures were devised for off-line editing and artifact rejection (Freeman & Schneider 1982), temporal filtering and decomposition (Freeman & Viana Di Prisco 1986b), spatial filtering and deconvolution (Freeman 1980; Freeman & Baird, in press), and multivariate statistical analysis of the results of measurement (Freeman & Grajski, in press; Grajski, Breiman, Viana Di Prisco & Freeman, in press). The procedures are reviewed elsewhere in detail (Freeman 1987b).

The measurement process consisted of curve-fitting of the 64 traces in each event. A set of 5 elementary waveforms or basis functions was identified as common to all 64 traces in varying degree. The sum of these 5 basis functions was fitted by regression to each trace, yielding 5 matrices of 64 amplitude values that incorporated 80% of the total variance of the event, as well as the matrix of residuals and the two matrices of the residues of high- and low-pass digital filtering, all expressed as root mean square amplitudes. Evaluation consisted of determining which of these 8 matrices best served (or served at all) to classify events correctly with respect to CSs and CRs. No data were discarded until they were tested in this way. Moreover, the coefficients of the basis functions were examined to determine whether they contained odorant-specific information.

The end results were unequivocal. The matrices of amplitude of the dominant basis function (the one containing the largest fraction of total power), and only these,

sufficed to classify events correctly. They did so at far above chance levels with respect to the two odorants in 4 of the 5 rabbits, who discriminated them behaviorally, but not the events recorded from the fifth rabbit, who failed to discriminate them (Freeman & Grajski, in press; Freeman & Viana Di Prisco 1986b; Grajski et al., in press).

An example of an event (unaveraged traces) is shown in Figure 3. The key property is that every trace had the same temporal waveform. Exceptions were due either to artifacts or to electrodes not placed on the bulb. The amplitude differed between channels so as to form a spatial pattern that (on the average) was relatively constant and easily identified with each animal. These amplitude patterns after familiarization remained constant unless and until odorant conditioning was undertaken. New patterns emerged only in association with reinforced odorants, not visual or auditory CSs or UCSs alone. They remained stable within sessions and across sessions provided the S-R contingencies were unchanged.

Multiple patterns emerged under discriminative conditioning. When a new odorant CS+ was introduced or when a previous CS+ was changed to a CS-, the entire set of spatial patterns appeared to change. The amount of change between stages that involved an altered S-R contingency, when measured as a fraction of the total between-session, within-stage variance, was relatively small (7%). The information in these stable spatial patterns that served to classify events correctly with respect to CSs and CRs was not localizable to subsets of channels. That is, the information density (as distinct from content) was spatially homogeneous, much as a letter-space on a printed page is of equal value whether it contains a letter, a punctuation mark, or no character at all.

3.2. The appearance of background activity. It is our belief that this is the first demonstration of the existence of sensory- and motor-specific information in the spatial dimensions of the EEG activity in any part of the cerebral cortex. The reason this has not been shown before is that

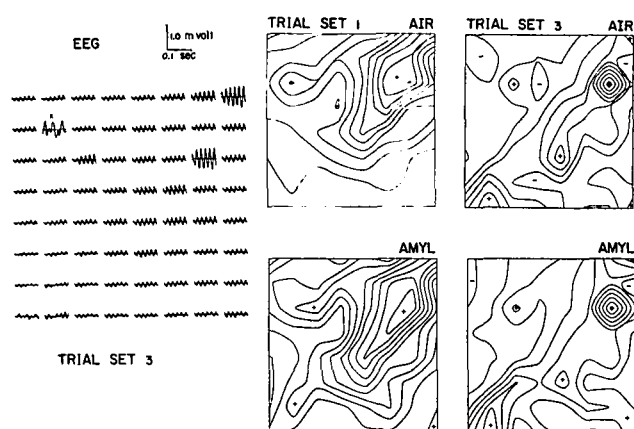


Figure 3. Left: A display of single unaveraged EEG traces is shown comprising a single odor burst among 10 bursts in a file from one trial set. The (x) marks an example of a bad channel record that was replaced during editing by an average of two adjacent records. Right: The root mean square amplitudes are compared for bursts without odor (above, "air") and with an odor (below, "amyl" acetate). There is a significant difference between the two patterns on the left but not the two on the right. From Freeman and Schneider (1982).

problems had to be solved at all levels of the project. These included practical problems such as array design and manufacture, surgical implantation, control and measurement of rabbit behavior, management of data flows on the order of 1.2 million bits per trial and several billion bits in each series of experiments, and basic theoretical problems in diverse fields including volume conductor analysis, statistical mechanics, nonequilibrium thermodynamics, nonlinear dynamics, and multivariate statistics applied to neural activity. The manufacture of arrays of electrodes, magnetic pick-ups, or optical probes and their preamplifiers merely opens the floodgates for the data. The difficult problems begin with the adaptation of recording to the conditions of normal, learned behavior, and with the rational design of algorithms for data reduction and refinement. Our methods happen to be the first that succeeded; there being no precedents, we have no other data with which to compare our results. Just as we have pioneered in their acquisition, we must now break new ground in attempting to understand what they tell us about brain function.

The elemental phenomenon that must be dealt with in olfaction, as in all of brain physiology, is the background activity manifested in the "spontaneous" EEG (Figure 2) and unit activity of neurons throughout the CNS. How does it arise, and what roles does it play? This activity is exceedingly robust; it survives all but the most drastic insults to cerebral tissue, such as near-lethal anesthesia, ischemia, or hypoxia. Perhaps the only reliable way to suppress it without killing the tissue is to isolate surgically small slabs of cortex (Burns 1958) by cutting neural connections while preserving the blood supply (and even then it may not be completely abolished). This procedure works for both the bulb and the prepyriform cortex (Freeman 1986), provided they are isolated from each other as well as from receptors and the rest of the brain. Under complete surgical transection of neural connections but with sufficient circulation for viability, each structure goes "silent" except when it is electrically or chemically stimulated. When perturbed and then left alone each structure generates a response and again falls silent. The responses to electrical impulse stimuli are observed through averaged evoked potentials (AEPs) and post-stimulus time histograms (PSTHs) of action potentials.

The state of a dynamic structure is said to be stable if the system returns to that state after perturbation. If the basal state is steady and nonoscillating, the system is said to be at an equilibrium. When the values of amplitude or energy are plotted on a graph, one against another, a response has the appearance of a curve or trajectory that ends at a point as the system goes to equilibrium. The same point is reached from many starting conditions under perturbation. Hence the point is said to represent an "attractor," and the set of starting conditions defines a "basin" for the attractor (Figure 4). When the system is placed by control of its input into the basin of an attractor, the system dynamics is said to be governed by the attractor.

When the stable equilibrium state of the bulb (Figure 2, bottom trace) or cortex is induced by deep anesthesia (Freeman 1986) or by cryogenic blockade of the axonal connections between the bulb and prepyriform cortex (Gray 1986) it is reversible. As recovery takes place the

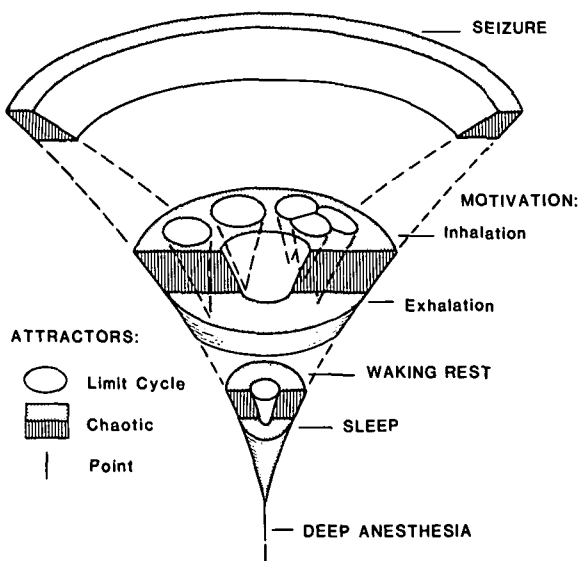


Figure 4. This vase-shaped structure is an attempt to portray a state space diagram for olfactory dynamics. The two horizontal dimensions constitute the axes for the amplitudes of activity of an excitatory subset and an inhibitory subset. The vertical axis serves to represent a bifurcation parameter, in this case the average level of driving input to the two subsets, consisting of input from centripetal activation of receptors and the input from centrifugal projection relating to arousal and motivation. The lowermost line represents an equilibrium or point attractor. Shaded areas represent a chaotic attractor, and the open circles represent limit cycle attractors. The activity for each stage is shown in Figure 2. A phase portrait derived from this diagram is shown in Figure 11. From Freeman (1987a).

background activity reappears; the system can be said to "bifurcate" or change to a new state, such that the point attractor is replaced by a point "repellor" (Figure 4). A repellor is manifested when attempts to quash or inhibit activity fail or succeed only transiently. The interconnected structures, the bulb and prepyriform cortex, cannot stay at equilibrium and must enter ceaseless activity, even if they are only connected to each other and not to the rest of the brain (Freeman 1987a). A bifurcation takes place when the system undergoes a major transition in its dynamics, equivalent to, for example, the transition from sleep to waking, or from normal to seizure activity. The governing equations are the same, but the solutions change radically. We say that the control of the system dynamics is shifted from a point attractor to a chaotic attractor. This simply means that the system falls into a condition of restless, but bounded, activity. It is stationary in the statistical sense, but its mathematical properties differ from those of "noise" (Grassberger & Procaccia 1983).

This background activity is statistically indistinguishable from what we call band-limited noise – that is, white noise passed through a band pass filter. We had known for years that the interval histograms of spike trains from single neurons conform to a Poisson process with a refractory period, so we had inferred that the background EEG was a local average of the dendritic potentials reflecting or governing the spike trains, a kind of "Brownian motion." In seeming confirmation of this view the correlation coefficient between pairs of traces fell with increasing

distance between their recording sites. From our recent studies we now know that this view was incorrect. The instantaneous frequency of bulbar EEG activity is always and everywhere the same, no matter how "noisy" the waveform may seem. The inverse relation of correlation with distance is due to small but systematic phase gradients extending over the entire bulb (Freeman & Baird, in press) and not to statistical independence of the samples. The commonality of waveform does not extend outside the bulb, but does extend over distances of several mm within it, much too far to be accounted for by volume conduction. The bulbar EEG is a global property that arises from dense feedback interactions within the bulb and yet is conditioned or made possible by extra-bulbar feedback interactions.

3.3. Evidence for chaos. An explanation of the neural mechanism of the background activity stemmed from our use of an assay, the Grassberger–Procaccia (1983) algorithm, to measure the degrees of freedom (the Hausdorff dimension) of a prolonged sample of the EEG from our animals at rest. Preliminary estimates ranged between 4 and 7 (Freeman 1987b), indicating that the activity reflected not "noise" but chaos (see note 1). This crucial distinction is analogous to the difference between the noise of a crowd at a ball game and the noise of a family dispute. Chaos is indistinguishable from random noise in appearance and in statistical properties, but it is deterministic and not stochastic (Garfinkel 1983; Rössler 1983). It has relatively small degrees of freedom; it can be turned on and off virtually instantaneously, as with a switch, through bifurcation (see sect. 3.2), unlike thermal noise, for example, which requires relatively slow heating and cooling. Chaos is controlled noise with precisely defined properties. Any system that needs random activity can get it more cheaply and reliably from a chaotic generator than from a noise source. Even the random number generators of digital computers are algorithms for chaos; given the same seed, sequences of random numbers are precisely replicated.

In order to replicate the EEGs of the olfactory system,




Figure 5. Impulse responses of the neural sets simulated for M (mitral unit activity), A of the AON (EEG), E of the PC (EEG), and G of the OB (EEG activity). The internal gains, k_{ee} , k_{ei} , and k_{ii} , are: OB (0.25, 1.50, 1.50, 1.80); AON (1.50, 1.50, 1.50, 1.80); PC (0.25, 1.40, 1.40, 1.80). The nonzero equilibria are not detectable with AEPs; the negative value for the PC is consistent with the silence of the PC after section of the LOT through the AON. From Freeman (1986).

we used sets of nonlinear ordinary differential equations that had already been used separately to model the bulb, anterior olfactory nucleus (AON), and prepyriform cortex (PC) with respect to their averaged evoked potentials (Figure 5). We coupled them into an interactive network (Figure 6). With proper settings of the feedback gains and distributed delays in accordance with our understanding of the anatomy and physiology of the larger system, the model yielded sustained chaotic activity that was statistically indistinguishable from the background EEG of resting animals (Figure 7). Under conditions of simulated receptor input the model generated "bursts" of oscillation that closely resembled those events seen in olfactory EEGs (Figure 8) during inhalation.

With some minor changes in gains between the bulb and AON the model system entered a degenerate state

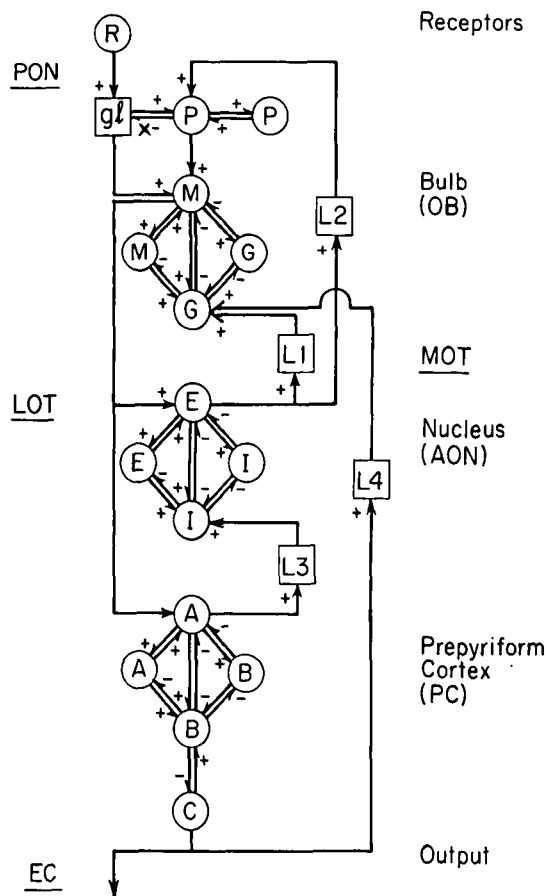


Figure 6. Flow diagram for the equation of the olfactory system. Each circle (except R) represents a second-order nonlinear differential equation (Freeman 1987a). Input from receptors (R) by the primary olfactory nerve (PON) is to periglomerular (P) and mitral (M) cells through the glomeruli (gl) subject to attenuation (x-), with connections to granule cells (G). Output by the lateral olfactory tract (LOT) is to the superficial pyramidal cells of the AON (E) and PC (A), each with inhibitory neurons respectively (I) and (B). Output of the PC is by deep pyramidal cells (C) into the external capsule (EC) and centrifugally to the AON and OB in the medial olfactory tract (MOT). The AON also feeds back to the granule cells (G) and the glomerular layer (P). Excitation is (+); inhibition is (-). Latencies (L1 to L4) are calculated from measurements of the conduction velocities and distances between structures. Each part is treated as a lumped system in this first approximation. Each path is assigned a gain - for example, $k_{MG} = k_{ee}$ in the OB, k_{ME} from the OB to the AON, and k_{EG} from the AON to the OB. From Freeman (1986).




Figure 7. Examples of chaotic background activity generated by the model, simulating bulbar unit activity (M) and the EEGs of the OB, AON, and PC. $Q_m = 5.0$, $k_{ME} = 1.5$, $k_{EG} = 0.67$, $k_{EP} = 1.0$, $k_{PM} = 0.1$, $k_{MA} = 1.0$, $k_{EA} = 1.5$, $k_{AI} = 1.0$, $k_{AP} = 1.0$

with a Hausdorff dimension near 2, manifesting a repetitive spike (Figure 2) that very closely resembled an epileptiform spike train that accompanied an electrically induced olfactory seizure (Figure 9). This phenomenon offered one means of studying the transition from a stable point attractor to a chaotic attractor (Babloyantz & Destexhe 1986) (Figure 4). We did this by increasing an excitatory gain connection in the model (k_{PM} in Figure 10, between sets P and M in Figure 6). This yielded the Ruelle-Takens-Newhouse route to chaos (Schuster 1984). The chaotic attractor of the "seizure" state of the model was a 2-torus; the chaotic attractor of a normal hyperchaotic background activity was much higher in dimension, and its geometric structure remained unknown. These results, which represent the first successful simulation of normal and abnormal EEG activity, and the experimental evidence supporting the mathematical model (Freeman 1987a) are reviewed elsewhere (Freeman 1986).

Given this broad picture of the dynamics of this neural system (Figures 2, 3, and 4) we can sketch a metaphorical picture of its multiple stable states in terms of a phase portrait. Each state is represented (Figure 11) by a surface in the two dimensions of the activity level of a representative local subset of excitatory neurons (left-




Figure 8. Left: A simulated burst induced by giving a surge of input at R similar to receptor input density during inhalation and exhalation lasting 0.2 sec. Right: Sustained input onto preexisting chaotic activity. From Freeman (1986).




Figure 9. Examples of 2-second time segments of EEGs recorded from a rat during a seizure, comparing these with the outputs of the model (see Figure 7). From Freeman (1987a).

right axis) and another of inhibitory neurons (axis in-out of the page). Vertical height in each place indicates the amount of energy in the active state of a point. An evoked potential would appear as a counterclockwise spiral trajectory; background activity would appear as a roughly circular squiggle around the base of the central projection. The equilibrium state of deep anesthesia is represented in the lowest plate at the bottom of a well. Its lowest point is the point attractor. The shift upward from one plate to the next depends on the degree of interaction within the system (the bifurcation parameter), which is subject to numerous parameters in the model and to various conditions in the brain relating to input and arousal. The sequence of bifurcation to the waking but unmotivated state is shown by the emergence of the central uplift, a point repeller, and the formation of a surrounding well that contains at its base the chaotic attractor. The state changes by which the central uplift occurs results in transfer of governance from a point attractor to a chaotic attractor (Figure 4).

Figure 11 indicates that the olfactory system and its corresponding model have a hierarchy of states. The basic neural dynamics and the equations are the same in all states but, depending on various neural conditions and model parameters, the systems behave differently (e.g., during waking, sleeping, bursts, interburst intervals, seizures, and so on). Both systems display the capacity for abrupt, dramatic, global jumps from one state to another. These are the bifurcations. These are analogous to phase

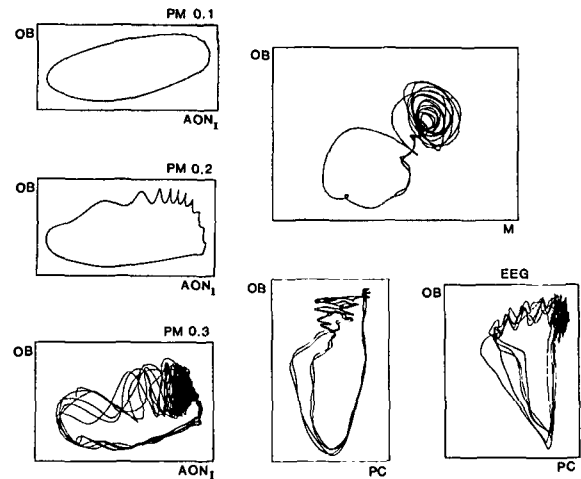


Figure 10. The traces at left show the spike train output (1.0 sec) of the model for 3 values of K_{PM} , showing a low-dimension limit cycle (above), a high-dimension limit cycle, and a chaotic attractor. (In one sense the limit cycle has only one dimension, that along its trajectory, but in another sense it exists in multiple dimensions, so that it never crosses itself.) Reconstruction of the chaotic attractor in 3 dimensions shows that it is a 2-torus without detectable orifices or folds. The upper-right frame shows a short segment (0.25 sec) from a different perspective. The lower-right frames compare the accessible OB and PC EEG traces during a seizure in a rat with the comparable output variables of the model. Although related, they are not identical. From Freeman (1987a).




Figure 11. A set of hypothetical phase portraits is constructed from the bifurcation diagram shown in Figure 4. Inhalation results in the emergence of the collection of learned limit cycle attractors, one of which may be selected by odorant input placing the system in its basin. Alternatively, the response may fall into the chaotic well. This appears to occur on about 10% of control inhalations and about 40% of the test odor inhalations after completion of training, as well as reliably with novel odorants (Freeman & Viana Di Prisco 1986). On exhalation the learned attractors vanish, so the system is freed to accept new input. At the top is the chaotic attractor of seizure; at the bottom is the point attractor of deep anesthesia. From Freeman (1987a).

transitions in physical systems: ice to water to steam, for example. The bifurcations occur in many forms and varieties, so a formal definition is difficult if not impossible to provide.

3.4. Roles of chaos in odor recognition. This configuration is retained under increasing motivation (as by food or water deprivation), resulting in higher amplitudes of background activity, but only during late exhalation. During late inhalation and early exhalation a surge of receptor input reaches the bulb, depolarizes the mitral cells, sensitizes the bulb, and induces an oscillatory burst. This is a bifurcation from a low-energy chaotic state to a high-energy state with a narrow temporal spectral distribution of its energy, suggesting that it is governed by a limit cycle attractor. Order emerges from chaos in two respects. First, a narrow spectral peak emerges, indicating high temporal coherence. Second, the local amplitudes of oscillation take on values that are reproducibly related to particular odorants serving as CSs. The values differ for different odors, indicating that multiple limit cycle attractors exist, one for each odorant an animal has learned to discriminate behaviorally, and each one leading to regular oscillation in a burst.

As hypothesized in Figure 11, these attractors are latent during late exhalation and in the absence of motivation. They reappear, all of them, with each inhalation under motivation and then vanish with exhalation. We postulate that the selection of an attractor upon inhalation is made by the presence of a CS odorant in the inhaled air or by the absence of an odorant, leading to the selection of an attractor corresponding to the background odor, the behavioral status quo. That is, the chemical stimulation of a particular set of receptors places the mechanism into a particular basin when the attractors emerge under bifurcation. The system is released into its basal state with exhalation, setting the stage for the processing of a new sample of information about an odor in the inhaled air.

The dominance of a chaotic attractor, perhaps in some sense closely related at all levels, is seen to extend from the low-level state of rest to the high-energy state of seizure. We conjecture that chaotic activity provides a way of exercising neurons that is guaranteed not to lead to cyclic entrainment or to spatially structured activity (Conrad 1986). It also allows rapid and unbiased access to every limit cycle attractor on every inhalation, so that the entire repertoire of learned discriminanda is available to the animal at all times for instantaneous access. There is no search through a memory store. Moreover, the chaotic well during inhalation provides a catch-basin for failure of the mechanism to converge to a known attractor, either because the sample is inadequate or because a novel or unfamiliar odor is present in the inhaled air. In either case a "disorderly" or chaotic burst results that is characterized by a relatively low peak frequency and a broad temporal spectrum reflecting excessive frequency modulation. Despite the spatial commonality of waveform, these bursts do not converge to a consistent spatial pattern of amplitude modulation, unless by repeated presentation under reinforcement a new CS and a new CR are formed, in which case a new limit cycle attractor emerges. In other words, the chaotic well provides an escape from all established attractors, so that an animal can classify an odorant as "novel" with no greater delay

than for the classification of any known sample, and it gains the freedom to maintain unstructured activity while building a new attractor.

In our view, then, chaos plays several crucial roles; the system is designed and built so as to ensure its own steady and controlled source of "noise" (i.e., chaos). Most remarkably, "signals" are not detected "in" the chaos because the mechanism turns the chaos "off" when it turns a signal "on." The immunity of EEGs to trauma shows that the mechanism is extremely stable, but not absolutely so. Petit mal type seizures (Figure 2) occur when the feedback control system is driven outside its normal range by excessive electrical stimulation and develops a dynamic asymmetry. This imbalance results in a pathological instability that carries the system temporarily into a degenerate and low-dimensional basin of chaotic activity; its pattern resembles the EEG spike activity seen during the early stage of recovery from "silence" under deep anesthesia. We believe that this common form of epilepsy manifests an "Achilles heel" of a common and widespread neural mechanism for the genesis and maintenance of various forms of chaos as the essential ground states of the perceptual apparatuses of the brain.

3.5. Learning and nerve cell assemblies. The neural mechanisms that underlie changes leading to the formation of a new limit cycle attractor have been described and discussed elsewhere in detail (Freeman 1975; 1979a-c; 1981; 1983b). Our model is based on studies of changes in the waveforms of averaged evoked potentials in the olfactory system when the electrical stimulus is used as a CS+ or CS-, and on replication of these waveforms by the impulse response solutions to differential equations simulating the dynamics of the bulb or cortex. Briefly, the excitatory neurons in each of these structures are synaptically linked by axon collaterals ending mainly on the cell bodies in bidirectional synapses (Willey 1973). When these neurons are co-activated pair-wise by a CS+ their joint synapses are strengthened in accordance with the Hebb rule (Viana Di Prisco 1984). The required reinforcement is mediated by norepinephrine, which is released into the bulb and cortex (and elsewhere) by the locus coeruleus (Gray 1986; Gray, Freeman & Skinner, 1986). Our models indicate that a modest increase of 25-40% in synaptic strength can increase the sensitivity of the bulb to a CS+ by 40,000-fold (Freeman 1979a; 1979b).

The linking together of a selected subset of neurons comprising perhaps 1-5% of the total by strengthened excitatory synapses constitutes the formation of a nerve cell assembly (NCA). Thereafter, excitation of any portion of it tends to disseminate into activating the whole of it. We imagine that each NCA exists as a filamentous network in the bulb resembling mold growing on a piece of bread. We hypothesize that the activation of some of the neurons of a specified NCA selects the basin of the attractor into which the bulbar mechanism converges on inhalation.

The key to understanding this switching device lies in an appreciation of the static nonlinearity that governs the behavior of neurons in an interactive mass. When left without input, neurons tend to fall below threshold and remain silent. Under maintained excitation they give steady output. Owing to the ionic mechanism of the

action potential there is a dynamic range near threshold in which the tendency to form an action potential increases exponentially with depolarization. Restorative forces released by an action potential serve to limit the rate of firing, but only after the fact, so to speak.

During exhalation, when receptor input is low, the bulbar neurons tend to fall to a low level of activity and sensitivity. During inhalation the surge of receptor input not only excites bulbar neurons, it augments exponentially their tendency to fire in response to input from receptors and from each other. Their strength of interaction increases dramatically over the entire bulb. At some point a threshold is reached in which the entire bulbar mechanism bifurcates from a low-energy chaotic state to a high-energy state. The NCA operates at the moment of choice when the surge of receptor input strongly forces the bulb far away from its rest state to some new activity pattern.

We view the bulb as operating in two modes. During late exhalation and early inhalation it is in a receiving or diastolic mode (Figures 2, 3, and 11). Intrinsic interaction strength is low. The activity of afferent axons is imposed on bulbar neurons, which are free to accept it and to adopt corresponding levels of firing. Both the temporal and spatial transfer functions are broadly tuned so as to accept information and maintain it by local firing (Freeman & Ahn 1976). This is the low-level chaotic state. On bifurcation the mechanism converts to the transmitting, or systolic, mode. Internal interaction goes to a high level. The temporal transfer function of the model changes to a sharp peak at the burst frequency, and the spatial transfer function changes to give a rapid fall-off in energy above zero c/mm . The bulbar neurons no longer respond to receptor input but instead to each other. The information carried by each neuron is disseminated over the entire bulb and is integrated by every neuron in the bulb. It is also sent out of the bulb to the cortex, where it undergoes further temporal and spatial integration. The integration is facilitated by the high temporal coherence of the oscillatory burst and by the occurrence of the burst center frequency in the optimal pass band of the prepyriform cortex viewed as a passive filter (Freeman 1975; Bressler 1987a,b). Feedback from the prepyriform cortex and AON to the bulb has the form of modulatory biases, because the conducting pathways have strongly dispersive delays that act as low pass filters and smooth the feedback activity. Upon reduction in receptor input during exhalation the system collapses back into low-level chaos and the diastolic mode.

3.6. Strong and weak points of our model. This view of olfactory discrimination arises from insights gained by inspection of the activity patterns revealed by these new data. It is consistent with most of what is known or believed about olfactory function from conventional electrophysiology, including the specificities of neuronal firing in response to odorant stimulation in anesthetized animals and the spatial patterns of selective 2-DOG (2-deoxyglucose) uptake in the glomerular layer (Lancet, Greer, Kauer & Shepherd 1982) on prolonged exposure of waking animals to odorants. It is also compatible with findings in olfactory psychophysics, particularly those relating to the relatively small number of odorants subject to absolute identification in the absence of prolonged

training or (in man) the use of verbal labels (Cain 1980). It also solves the problem of neuroanatomical interfacing by the bulb between the receptors and the primary olfactory cortex as follows.

The input path to the bulb, the primary olfactory nerve (PON), has a certain spatial organization that is imposed by ontogenetic development and by functional needs to be met in getting receptor input into the glomeruli in the face of lifelong replacement of the primary receptors (every 120 days, on the average). The output path, the lateral olfactory tract (LOT), has its own constraints in its ontogeny and in the need to service an array of targets ranging from the AON and tubercle to the amygdaloid nucleus and hippocampal rudiments. By our hypothesis, within a few msec following bifurcation all information that is fed into the bulb during its "diastole" (the inter-burst period) is spread and mixed uniformly through the bulb during its "systole" (the burst). Each fraction of the bulbar output, perhaps on the order of 20%, irrespective of which part of the bulb it comes from, suffices to convey with adequate resolution all that the bulb has to say. Hence there need be no coordination or sharing of constraints in the developmental construction of the input and output paths, particularly with respect to their topographic organizations.

Several challenges and uncertainties exist for the physiology of our model. One of the key features by which it must stand or fall is its requirement that the interneurons in the outer layer of the bulb (Figures 1 and 6), the periglomerular cells, must be excitatory to each other and to mitral cells (Freeman 1987a). Substantial but indirect experimental evidence has been adduced in support of this requirement (Martinez & Freeman 1984) as well as against it (Shepherd 1972). The cells in question are mixed populations of cells secreting GABA, dopamine, and one or more neuropeptides. Conventional wisdom has it that small GABA-ergic neurons are inhibitory. This appears to be valid for the deep-lying granule cells. Recent studies in the hippocampus have shown that GABA is hyperpolarizing when applied to the basal dendrites of pyramidal cells but depolarizing when applied to the apical dendrites (Misgeld, Deisz, Dodt & Lux 1986), suggesting that a chloride gradient along the apical dendritic shafts might reverse the sign of action of GABA between the two parts of the pyramidal cells (Newberry & Nicoll 1985). Were this or an equivalent mechanism to hold for the mitral cells in the bulb, an important prediction by our model would be confirmed.

Another key property of our model is the requirement of mutually inhibitory feedback among inhibitory interneurons in the bulb, AON, and cortex. No direct demonstration that this does or does not exist has yet been devised. Evidence for chemical and electrical synapses between granule cells has been sought but not found. The possibility exists that the stellate cells of Golgi, Cajal, and Blanes, which are thought to be GABA-ergic, inhibit granule cells through their widely distributed axons, and might receive inhibitory input from them.

A third weakness concerns the requirement for mutually excitatory connections among mitral cells in the bulb and among superficial pyramidal cells in the prepyriform cortex, which are modified under learning. The evidence that these requirements hold comes largely from recordings of field potentials and is therefore indirect. Studies of

the predicted synaptic changes and their kinetics under modulatory neurochemical agents may be crucial for the support of our model. However, we emphasize that the jury is still out on these questions, that an answer of a particular kind is required for each by our model, and that some other answers can falsify it. We therefore have a brain theory that can be tested, elaborated, or negated by physiological experiments; it is not merely computational.

In its mathematical structure our model is still in its infancy. We have some experience with a distributed system of coupled equations in bifurcation between equilibrium and limit cycle states (Freeman 1979c), but our chaotic generator is a lumped model using ordinary differential equations. In its psychological dimension our model is extremely limited, being competent to simulate only preattentive cognition (Freeman 1983a; Julesz 1984) and the instantaneous apperception of a stimulus, and not attentive inspection or sequential analysis. The feasibility of extending these ideas and experimental methods to neocortical systems is under exploration; evidence has been found that the visual cortex in a rhesus monkey operates according to the same basic neural dynamics as the olfactory bulb (Freeman & van Dijk, submitted). Most important, no claim for firm and substantial understanding of large-scale neural circuitry can be advanced until the mathematical theorists of distributed dissipative systems have caught up with experimentalists, or until engineers have built hardware models based on our equations and determined whether they behave the way parts of brains do. We are pleased to present something new to think about.

4. Philosophical aspects

4.1. Neural dynamics and the digital computer. Our present hypothesis is that odor discrimination and recognition depend on self-organizing neural processes in the olfactory bulb. The process that we label the "expectation" of an odor is realized in the formation of strengthened connections in a network of neurons constituting the NCA. This assembly, whose role is to amplify and stereotype the small input received on any given inhalation, produces a disseminated but low-density activity pattern in response to the stimulus, and then provides the crucial mechanism for mediating the emergence of an odor-specific activity pattern in a process of bifurcation. With this state change the entire olfactory bulb, rather than the limited number of nerve cells comprising the NCA, is engaged by a process of global integration to produce a stereotypic activity pattern mediated by the NCA but going far beyond it. Thus, when placed in a learned input domain, the neural system has a tendency to generate a qualitatively distinctive form of ordered behavior that emerges from the chaotic background state.

Several important lessons concerning recent explanatory models in cognitive science can be drawn from our research. First, our model, based on self-organizing neural dynamics, makes it desirable to reevaluate the adequacy of the explanatory models based on digital and analog computers that have until recently provided the most influential metaphor in cognitive science. According to this metaphor, the behavior of a system is caused by

the formal manipulation of bits of data (symbols) according to rules and operations specified by programs designed for a given task or tasks. The metaphor involves a distinction between system hardware and software, the functioning of a central processor that operates on the data and drives the system, and a memory housed in a separate space. Several factors have contributed to the decline of this metaphor, among them the evidence that implementations of it fail to produce behaviors in which animals and humans excel (Dreyfus & Dreyfus 1986) and the emergence of alternatives in the form of "connectionist" models.

Our data indicate that what takes place in the olfactory system does not resemble the processes responsible for generating behavior in the classical computer paradigm. In the olfactory bulb, learning consists in the selective strengthening of excitatory connections among the neurons leading to the constitution of an NCA and to the possibility of bifurcation to a global activity state manifesting an attractor. Learning takes place during the first 2 seconds following odor CS+ and UCS presentation with the release of norepinephrine in the bulb and elsewhere. Memory for an odor consists in the set of strengthened excitatory connections of the NCA, which, when activated under stimulus input, possesses the tendency to produce a global activity pattern characteristic of a given odor. These are not the types of mechanisms used by digital or analog computers. No program-specified rule or operation is brought to bear on input to the olfactory system. The component neurons generate their own ordered response to stimuli; they are self-organizing. There is no central processor, and learning and memory are functions distributed throughout the neural network.

The process of odor recognition and discrimination can be conceived in terms of dynamic interactions at the level of the neural mass without appeal to symbols. There is preliminary evidence from anatomical and EEG studies indicating that this distributed model can be generalized for neural dynamics throughout the cortex (Freeman & Skarda 1985; Freeman & van Dijk, submitted). This means that the classical computer analogy may be unsuitable to explain the neural bases of behavior. This does not mean that digital computer models are to be discarded. Von Neuman machines have successfully produced some interesting classes of behavior, and to date psychological models seem to lend themselves more simply to formulations stated in terms of symbols and their formal manipulation by rules. What we wish to point out here is that brains do not use the same principles as the digital computer to produce behavior. This information may help neurophysiologists in framing hypotheses for further research. Rather than viewing brain function along the lines suggested by the classical computer paradigm – as a rule-driven and controlled system solving problems, completing patterns, and forming hypotheses by manipulating symbols – neural dynamics suggest that the brain should be viewed as a self-organized process of adaptive interaction with the environment.

4.2. Neural dynamics and connectionist models. Our model supports the line of research pursued by proponents of connectionist or parallel distributed processing (PDP) models in cognitive science (Baird, in press; Feldman & Ballard 1982; Hinton 1985; Hinton & Anderson

1981; Hopfield 1982; Kohonen 1984; Rumelhart, McClelland & PDP Research Group 1986). Although the models that fall under the rubric of connectionism are not identical, they do share a number of basic characteristics (Feldman & Ballard 1982). Each involves a processing system consisting of a densely interconnected network of units that interact with one another by sending and receiving signals modulated by the weights associated with the connections between the units. Processing is distributed throughout the system. The units may be organized into layers, and each layer sends to and receives signals from other layers composed of densely interconnected units. The state of each layer results from a synthesis of the states of other layers from which it receives input.

What takes place in the brain may resemble the dynamic processes of self-organization used by these models. Our neural model and the connectionist models converge in several respects. Both rely on parallel, distributed processes among highly interconnected units in interacting networks to produce behavior; both emphasize a self-organized or bottom-up, rather than a rule-driven or top-down, explanatory approach; and both rely heavily on organized feedback among components within the system.

The convergence of our model with connectionist models is instructive. Equally striking are the dissimilarities (Baird, in press). Comparing the models shows that the study of brain dynamics provides essential information about the physical processes responsible for behavior that is not available from current engineering research alone. Our data show that neural dynamics exhibit features not found in connectionist models, features that we hypothesize are essential for odor recognition and discrimination. Modifying the connectionist models along these lines could yield more flexible systems capable of operating successfully in a more realistic environment.

4.3. Feedback of multiple kinds. The first point of difference between our model and connectionism concerns the process of feedback. Neural masses possess (and their collective dynamic behavior is determined by) dense local feedback among the neural units comprising the bulb and within it the multiple existing NCAs. This property is essential for the complicated dynamical processes of neural interaction needed for state changes and for the chaotic and limit cycle behaviors discussed above. Without locally dense feedback formed by the dendritic plexus that provides for a continuum of local interactions in a spatially distributed manner, the dynamic processes responsible for odor recognition and discrimination could not take place.

There are two points to make regarding feedback in connectionist models. First, approaches of the perceptron class (Hinton 1985; Rosenblatt 1962; Rumelhart et al. 1986) do not realize this kind of feedback in their models, even to the extent circumscribed by limitations on the hardware. Some models of this class do involve feedback from one layer of units to another layer (Hinton 1985; Rumelhart et al. 1986). In these models the activity of one layer (B) is fed back to a previous layer (A), thereby modifying the weights of the units in layer A. In contrast to these models the process of "backward propagation" in the brain imposes long delays, temporal dispersions, and

spatial divergences that do not hold for local feedback (Figure 5, L1-L4). Feedback between layers (e.g., between bulb and prepyriform cortex) is not equivalent to the short-latency, focused feedback taking place in the neuropil. There each node has surrounding plexuses of connections which concurrently excite and inhibit by recursive actions. Second, most connectionist models (e.g., Anderson, Silverstein, Ritz & Jones 1977; Hopfield 1982; Kohonen 1984) have excitatory feedback dominantly or exclusively; the role of inhibition has received scant attention. Models based on symmetric matrices of connection weights cannot simulate neural functions because of the existence in the nervous system of the mix of positive and negative feedback. There are exceptions in the connectionist literature; the models developed by Grossberg (1980) feature the local inhibitory feedback found in the neural mass. But the types of dynamics and of connectivity in those models still do not approach those occurring in the bulb and its NCAs, and we doubt that they can produce the global behaviors that characterize its neural dynamics.

4.4. Roles of chaos. A second, related point of difference between neural dynamics and connectionism involves the nature of dynamic behavior exhibited, on the one hand, by neural masses, and on the other hand, by connectionist networks. Our data support the hypothesis that neural dynamics are heavily dependent on chaotic activity. We have suggested that without chaotic behavior the neural system cannot add a new odor to its repertoire of learned odors. Chaos provides the system with a deterministic "I don't know" state within which new activity patterns can be generated, as shown by what happens when the system encounters a previously unknown odor. If the odor occurs without reinforcement, habituation takes place; thereafter, the neural system exhibits patterned activity that we have identified as the control state for the status quo. With reinforcement, however, a completely different process occurs. If the odor is novel and the system does not already have a global activity pattern corresponding to the odor, then instead of producing one of its previously learned activity patterns, the system falls into a high-level chaotic state rather than into the basin for the background odor. This "chaotic well" enables the system to avoid all of its previously learned activity patterns and to produce a new one.

In the neural system, we postulate that the process of state change leading to the unstructured chaotic domain is essential for preventing convergence to previously learned patterns, and hence for the emergence of new patterned activity in the bulb and the corresponding ability of the animal to recognize new odors. In the olfactory system the chaotic background state provides the system with continued open-endedness and readiness to respond to completely novel as well as to familiar input, without the requirement for an exhaustive memory search.

Connectionist models can certainly be modified to produce chaotic and oscillatory behavior, but current theorists have not included these behaviors in their models, nor have they adequately explored the potential benefits of doing so. One reason for this may be that we all lack the appropriate mathematical tools to implement

these behaviors at the spatial computational level. Another is that engineers have traditionally viewed oscillatory and chaotic behaviors as undesirable and something to be eliminated (Garfinkel 1983).

The connectionist model reviewed by Hopfield and Tank (1986) is instructive in this regard. This model captures the dynamics of a system at a point in time after the bifurcation included in our model from diastole to systole has taken place. It is pictured in their phase portrait by a set of point attractors. Input places the system into the basin of one or another of these point attractors, to which the system then converges. There are two problems with this model from our perspective. First, the neural system does not exhibit behavior that can be modeled with point attractors, except under deep anesthesia or death. Convergence to a point attractor amounts to "death" for the system. In the Hopfield and Tank model, after the system converges to a point attractor there is no intrinsic mechanism by which the system can escape from it. An obvious solution is to turn the system off and then to reset it so that it is free to converge again to another point attractor. This is like using a muzzle-loader instead of a machine gun. Second, their connectionist model lacks an intrinsic mechanism like the chaotic well in our model that enables the neural system to add new odors to its repertoire. Without such a mechanism the system cannot avoid reproducing previously learned activity patterns and can only converge to behavior it has already learned. The neural system does not have this problem; chaotic mechanisms enable the neural network to learn new behaviors.

4.5. Pattern completion versus destabilization. A third difference between neural dynamics and connectionism concerns the general conceptual framework in which the two models are explained. Connectionist models are sometimes understood as pattern completion devices, in which, for example, when the receptor units are given part of a pattern as input, the complete pattern can be reconstructed by interactions among appropriately weighted units comprising the network.

The neural system we have described is not best thought of as a pattern completion device, although it may do that (Freeman 1983b). The problem is epistemological; we do not know what a completed pattern is (so convergence to it cannot be ascertained as in an error correction device), nor, we suspect, does the brain. We postulate that an NCA is activated wholly by input to any of its neural members, but we have no measure or observation of what the NCA looks like or how completely it is activated. The output of the system does not consist of the "completed" pattern of the NCA but of the entire bulb governed by an attractor. This global state tends to recur within certain "clusters" of spatial patterns, possibly expressed as vectors in some high-dimensional space, but no two are identical, and there is no expression for a boundary, such as the outline of a letter that is to be filled in. Most generally, these neural activity patterns are generated from within. Whatever "meaning" they have is embedded in the self-organized matrix of the entire brain. We have no way of knowing what constitutes a "completed" pattern or how to distinguish it from an "incomplete" one, either in terms of neural activity patterns or the mental life of an animal, presuming it

exists. The pattern-completion concept is realizable only in terms of ideographs or conventional signs and symbols, and if we reject these, as we have for neurophysiology, then the concept too must go.

We also think that the term "pattern" in the expressions "pattern completion" and "neural activity pattern" has very different connotations and different implications for our understanding of system dynamics. The term "pattern completion" describes a process in which a circumscribed structure can be generated as output from input that provides information about only part of the structure. Generally, the process depends on a prior "optimal" presentation of the pattern to adjust the weights among units comprising the network. The neural system works differently. It cannot depend on optimal input in its first (or, in fact, any) encounter with an object against which to compare or judge subsequent input. In the neural system, chaos is the rule, and the patterned activity to which the system converges following each state change is never twice the same, so again the notion of pattern completion loses its meaning.

We think that the notion of "destabilization" provides a better description of the essentials of neural functioning than the concept of pattern completion. In an alert, motivated animal, input destabilizes the system, leading to further destabilization and a bifurcation to a new form of patterned activity. We hypothesize that convergence to an attractor in one system (e.g., the olfactory bulb) in turn destabilizes other systems (e.g., the motor system), leading to further state changes and ultimately to manipulation of and action within the environment. Our research leads us to postulate that behavior can best be modeled as a sequence of ordered, stable states in an evolutionary trajectory (Freeman & Skarda 1985). Input to the system continually destabilizes the present stable state and necessitates convergence to a new form of behavior.

4.6. The sensory/motor loop. This raises a fourth issue. The fundamental character of behavior is adaptive interaction in the world (Churchland 1986; Skarda 1986). Feedback from the consequences of behavior modifies the system and projects it into a higher order of stability. In the nervous system, each change in the dynamic structure of the system, which in our mathematical model requires a new solution with its own trajectory, occurs in a sequence of state changes. The convergence to a patterned activity state, which marks the end state of this process, is externally manifested in some physical state (e.g., chewing) or in some anatomical structure (Ermentrout, Campbell & Oster 1986). What is important here is that this state has a musculoskeletal pattern that constitutes both input to and output from the nervous system. The global activity pattern we record is the result of the destabilizing effects of receptor input to the system, but it is likewise the cause of motor output (e.g., licking) that causes further sensory input and manipulation of the environment, as well as being the result of previous motor activity (e.g., sniffing). These global patterns of the nervous system are at all times locked into both sensory and motor patterns of input and output.

We know that the neural system accomplishes this, but our model clearly does not contain a description of the mechanism by which this interaction is achieved. There

are existing prototypes that can be drawn upon by theorists in the work of Walter (1953), Ashby (1952), and Grossberg (1980) of systems with perceptual processes that interact with the environment via motor functions. Further interdisciplinary research along these lines is required before the mechanisms responsible for adaptive interaction will be understood.

Our data lead us to the view that the neural processes of self-organization in the olfactory bulb are quite selective. The olfactory system does not respond to each odor presented to it by producing a corresponding activity pattern. Neural dynamics and the formation of patterned activity that can be correlated with a specific odor are a function of "motivation." What we have labeled motivation can also be characterized as a complex process whereby the organism predictively controls and maintains itself in the optimal condition given the circumstances in which it exists and acts. These global objectives constrain neural dynamics in two ways: They limit the possible range of patterned neural behaviors and they mediate interaction among various neural subsystems, such as that between the brain stem including reticular formation, the locus coeruleus, and the olfactory system relating to arousal, attention motivation, learning, and so on.

Nonbiological self-organizing systems are different. Their behavior is not constrained by either of the global constraints operating in the neural system (Ashby 1952). Storms, for example, are self-organized phenomena that can be mathematically modeled using the same principles we use to model neural dynamics. A storm takes in and gives out energy, moves across a random path buffeted by external forces, and finally dissipates when it has depleted its energy sources. Storms, however, do not exhibit adaptive responses: The system dynamics of a vortex are not constrained by a global demand for preservation of the system, and the system does not incorporate information about its environment. The storm may, for example, move toward land, but it does not do so under the constraint to survive as a unity.

The distinction we have drawn between brains and nonbiological forms of self-organization does not guarantee that brain dynamics will always exhibit the self-promoting constraint just outlined. Sometimes brains produce behaviors that resemble the system dynamics characteristic of weather patterns; we identify some "neural storms" as seizures (Freeman 1986). The presence, however, of a self-organized neural process that is not self-promoting disrupts normal functioning at all levels. These otherwise common and efficient nonbiological forms of self-organization take on a pathological character when they occur in the brain. We propose that they are identified as pathological because they violate the global constraints for self-promotion and adaptive control characteristic of normal brain functioning. Thus, a difference between biological and nonbiological forms of self-organization shows up at the level of the neural assembly long before there is any reason to refer to "consciousness" or "beliefs." Self-preservation plays a central role in biological self-organized systems, and processes that do not possess this feature may be selected against during evolution.

The constraints on self-organization operative in the nervous system are not present in all biological systems.

For example, plants exhibit adaptive behavior indicating that their systems are governed by the constraint for self-promotion over time. Eucalyptus trees influence (modify) their environment by inhibiting the growth of other tree species and shrubs in their vicinity, and they promote their own rainfall from coastal fog. These are clear examples of self-promoting control of the environment, a feature that sets apart biological self-organized systems from nonbiological ones. But brains introduce a further constraint, not found in other biological forms. The key property of brain dynamics, we suggest, is control of body movement in space for the self-promoting purposes of search, attack, ingestion, escape, and reproduction. Plants have no brains. This is why we claim that there can be no adequate explanation of brain function without consideration of sensation in conjunction with movement. Nervous system dynamics is a self-organized process constrained by the requirement that the system anticipate and incorporate the immediate consequences of its own output within the larger constraints of regulating its well-being and the long-term optimization of its chances for survival. This is subsumed in J. J. Gibson's (1979) theory of "affordances." We are a long way yet from understanding how brains accomplish this.

ACKNOWLEDGMENT

Supported by grant MH06686 from the National Institute of Mental Health.

NOTE

1. "Chaos" in the oldest sense means the formless void from which order springs. The term is now commonly applied to disorderly, unpredictable happenings that give an observer no sense of regularity. In the technical sense used here it describes a kind of activity that appears to be random or stochastic by every standard statistical test, 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. It is pseudorandom noise that can be reproduced with high precision if the initial conditions are identical on repeated runs, but which is unpredictable if new initial conditions are used. In contrast to noise, chaos has fewer degrees of freedom and is said to be low-dimensional. Chaos exists in many forms and degrees; Rössler (1983) has formulated an instructive hierarchy of equations to exemplify types of chaotic activity that will be of great interest for neural theorists. Introductory texts are by Schuster (1984) and Abraham and Shaw (1985).

Open Peer Commentary

Commentaries submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Chaotic dynamics in brain activity

A. Babloyantz

Faculté des Sciences, Université Libre de Bruxelles, Campus Plaine CP 231, 1050 Brussels, Belgium

Skarda & Freeman (S&F) attempt to extend the recent advances in the analyses of nonlinear dynamical systems to the study of the olfactory bulb.

The suggestion that brain activity has chaotic deterministic dynamics is not new and has already been proposed by several authors (Kaczmarek & Babloyantz 1977; Nicolis 1985a; Roschke & Basar, in press). The degree of chaos in various stages of sleep and wakefulness has been evaluated from human EEG recordings (Babloyantz & Destexhe, in press; Babloyantz, Nicolis & Salazar 1985; Layne, Mayer-Kress & Holzfuss 1986; Rapp, Zimmerman, Albano, Deguzman & Greenbaun 1985). A low-dimensional deterministic chaos was found in an episode of petit mal seizure (Babloyantz & Destexhe 1986).

Reference to the above-cited material might have strengthened the main point of S&F's target article, namely, that brain activity in animals as well as in man conforms to deterministic dynamics of a chaotic nature.

In order to describe the various concepts of nonlinear dynamics without the help of mathematics, several misleading statements are introduced by S&F. "Chaos is controlled noise with precisely defined properties" (Sect. 3.3, para. 1) is an example. To understand S&F's paper, the reader must refer to more technical publications.

In spite of these remarks, S&F's combination of multiple electrode recordings and dimensional analyses is a very promising method for analyzing brain activity. Such an approach sheds light on some aspects of cerebral dynamics not accessible by other methods.

Chaos, symbols, and connectionism

John A. Barnden

Computer Science Department, Indiana University, Bloomington, Ind. 47405

As I have no doubt that the study of chaotic behavior in neural networks is an interesting and fruitful line of research, I shall confine my comments largely to some philosophical claims Skarda & Freeman (S&F) make and to the relationship of their model to connectionism.

Governing metaphors should be kept constantly under review as a matter of principle. But I am not convinced that S&F's findings and arguments constitute a serious threat to the "digital computer metaphor," or, more precisely, to views of the brain as a symbol-manipulation device. There is, first, the danger of extrapolating from findings and theories about low-level sensory mechanisms to high-level cognition. S&F are sensitive to the danger and freely admit that their model is extremely limited psychologically (see end of Sect. 3.6), but they would nevertheless like us to allow the extrapolation. More important, however, S&F do not give us any *new* reason to be worried about viewing higher-level cognitive processes as based on symbol manipulation. As far as I know, the "symbol manipulationists" have in any case always presumed that low levels of perception are at least largely based on specialized mechanisms that are probably not to be regarded profitably as manipulating symbols in any conventional sense. [The fact that AI (artificial intelligence) researchers and others simulate such mechanisms on digital computers is of course only weakly relevant here.] To show, therefore, that the olfactory bulb is best described as operating in a way foreign to symbol manipulation is not to push the backs of the symbol manipulationists any nearer to the wall.

The question of what the rest of the brain does with the output

of the olfactory bulb (and directly connected brain centers) is significant. S&F themselves come near to suggesting that the bulb produces symbols when they say the inhalation of a learned odor pushes the bulb into a qualitatively distinctive, stereotypic state of activity. What is to stop us regarding these patterns as symbols? In what way is the idea that the rest of the brain uses these patterns in a symbol-manipulation style rendered implausible? I am not arguing that the rest of the brain does so use them, but only wondering what light is thrown on the issue by the S&F model.

S&F might reply that the activity pattern resulting from a specific learned odor varies somewhat in response to environmental context and preexisting internal state, and therefore cannot be regarded as a symbol. This point has some force, but there appears to be nothing to stop me from retreating slightly and saying that the rest of the brain proceeds to *extract* a symbol – corresponding to some invariant part or aspect of the pattern. (That such a part or aspect exists is surely at the basis of S&F's model.) Also, no retreat at all might be necessary if we allowed symbols to embody a certain amount of "fuzz." This would depart from the conventional view of symbols in artificial intelligence and cognitive science, but it is not clear that the unfuzziness of symbols is crucial to those fields, even if most researchers in those fields think it is. The most crucial aspect of the symbol-manipulation view seems to me to be the ability to form complex structures out of basic symbols, to analyze such structures, to compare symbols, and to associate symbols with symbols and other entities. None of these abilities requires unfuzziness of symbols in principle [Nelson Goodman's (1968) views on notational systems notwithstanding].

S&F would do well to be more careful about nomenclature when making their philosophical claims. The metaphor that their attack is directed at is surely the symbol-manipulation metaphor, not a metaphor of the brain as a digital computer as such, since it is clear that the brain is not like a computer at a low level of description. Now, at the electronic level of description a computer does not operate by symbol manipulation any more than a neural net does at the neurophysiological level of description, so that any terminology that confuses levels is likely to be misleading. When, for instance, S&F say in Sect. 4.1 that the process of odor recognition and discrimination can be conceived in terms of dynamic interactions at the level of the neural mass without appeal to symbols, we might well respond that the behavior of a program running on a computer can be conceived in terms of dynamic interactions at the level of the electronic mass without appeal to symbols. That this response would not (I take it) get at the heart of what the authors are saying would be their own fault, to put it abruptly.

While we are on the subject of levels, I dispute the implication in Sect. 4.2 that "self-organized" is correlated with "bottom-up" or that "rule-driven" is antithetical to "self-organized." A rule-driven system can be self-organized at the level of rules (since rules can modify themselves and other rules), and a top-down decomposition of a system can involve elements of self-organization at any level. Actually, it is not clear that either connectionists or S&F are adopting a bottom-up approach. To be sure, they are suggesting particular low-level mechanisms to explain particular high-level behaviors, but that does not make the approaches bottom-up. It is more that they are adopting a top-down approach different from those taken by certain other researchers.

I am a little puzzled at the claimed divergence from connectionism with respect to types of feedback (Sect. 4.3). There seems to be nothing in the spirit of connectionism that disallows "locally dense feedback." Also, inhibitory feedback, which is claimed by S&F to have been given scant attention in connectionism, has played a very significant role in connectionist thinking for some time. One need only look, for instance, at the model of McClelland and Rumelhart (1981) and at the importance given to lateral inhibition in the Kohonen (1984) book

cited by S&F. On the other hand, I do agree that connectionists would do well to look more closely at the transmission delays and temporal dispersion effects on connections (whether or not they are feedback connections). I think it is really best to regard S&F's reliance on chaos and certain feedback effects as constituting a (most intriguing) extension of present-day connectionism rather than as diverging from it.

Finally, I would be interested to know what happens in the olfactory bulb and in S&F's model when several individually learned odors are presented simultaneously. Can a spurious output result, by virtue of the combined odors pushing the system into an activity state corresponding to another learned odor? Does this, if it happens, have any correlation to observed behavior? What new light is thrown on whether the output patterns can be usefully viewed as taking part in symbol manipulation?

Spatial analysis of brain function: Not the first

Robert M. Boynton

Department of Psychology, University of California at San Diego, La Jolla, Calif. 92093

Skarda & Freeman (S&F) believe that their work is "the first demonstration of the existence of sensory- and motor-specific information in the spatial dimensions of the EEG activity in any part of the cortex" (Sect. 3.2, para. 1). Not so: Their attention should be directed to *Toposcopic Studies of Learning*, a book by Donald W. DeMott (1970). In addition to reporting his own work on toposcopy (study of the cortex in two dimensions of space), DeMott reviews the history of the subject, stating that two earlier reviews had been published previously, one by A. Rémond in 1955, the other by Livanov and Anan'yev in 1961 (which I will not cite directly, because I have not seen them). He states that the first toposcopic experiments that produced useful data were published by Lilly and Cherry (1951; 1954) using a 25-channel apparatus.

DeMott studied a variety of learning problems in the monkey while trying to record simultaneously from as many as 400 electrode positions in the brain. He describes both AC and DC changes in the recorded potentials and relates his results to such phenomena as dominant focus, contingent negative variation, hemispheric dominance, and localization of function. For example, in his "one tone, one-string problem," a monkey received a grape reward for pulling the string, contingent upon the presence of the tone. Five Cebus and four Saimiri were studied over several sessions. On some trials, toposcopic patterns, in the form of an array of lights whose intensities were proportional to brain potentials, were recorded with a high-speed camera of original design at 250 frames per second. DeMott discerned a distinctive pattern of electrical activity associated with the first behavioral signs of learning, one that straddled the parieto-occipital sulcus. Such activity was never otherwise observed, even during analogous visual learning studies. He refers to this activity as an apparent "lateral movement of activity in the region of the focus . . . sharply limited, as if by an invisible fence around the critical area" (p. 93).

DeMott's book also includes detailed discussion of the design and manufacture of electrode arrays and of the problems encountered with respect to surgical implantation as well as the formidable problems of data analysis in what was, for him – given his limited resources – the precomputer era.

S&F also state, "there being no precedents, we have no other data with which to compare our results. Just as we have pioneered in their acquisition, we must now break new ground in attempting to understand what they tell us about brain function" (Sect. 3.2, para. 1). As I hope the foregoing will attest, these

matters are what DeMott's monograph is largely about. Otherwise, it is a personal history (I find it a sad one) of what can happen to those who arrive before their time and choose to go it alone. With uncommon frankness, he tells of his long struggle with journal editors and grant reviewers, who eventually put an end to his research career on August 31, 1968, ten years after the toposcopic project had begun. His efforts deserve to be remembered.

Such carping out of the way, I will conclude by saying that I otherwise enjoyed the paper by S&F. Better than most of us, they have utterly banished the homunculus, or "green man," from their thinking and have called attention to the fundamental weaknesses of the simple-minded brain-computer analogies. Yet I find S&F's records difficult to interpret, just as reviewers found DeMott's. Only time will tell whether chaos is in fact the route to making sense of the world.

Can brains make psychological sense of neurological data?

Robert Brown

Department of Psychology, University of Exeter, Exeter EX4 4QG, England

Churchland (1980) distinguishes two varieties of scepticism concerning the usefulness of brain research for our understanding of how the mind-brain works: "boggled" and "principled" scepticism. Presumably this distinction is only a makeshift one, and scepticism merely an imprecise "folk-psychological" notion, but since the terms have not yet been eliminated from our psychological vocabulary I will assume that they are still meaningful. Skarda & Freeman's (S&F's) target article, for all its good intentions, increases my scepticism (of both kinds).

Being blinded by science could well be a function of one's own intellectual eyesight. No one expects the general theory of relativity to be easily assimilable in comic-book format, but an argument is still expected to meet the criteria of clarity and intelligibility. In their account of the collection, analysis, and interpretation of data S&F mystify and intimidate, albeit unwittingly. It must be a small and specialised community indeed that can follow each of the technical and mathematical steps with a truly critical eye. This is not a trivial criticism. It is claimed that the model can be tested, elaborated, negated, or falsified. We know that falsification (or potential falsification) is not the cut-and-dried procedure it was 'once thought; given such sheer complexity, what would it really take to falsify this model? There could be many an inferential slip 'twixt sniff and sip, and how many would honestly be the wiser?

However, boggled scepticism is not a serious complaint; it can be cured in this case (in principle) by getting down to the hard work of understanding volume conductor analysis and the Ruelle-Takens-Newhouse route to chaos (Schuster 1984). Clashes of principle are clearly more serious, occasionally terminal, disorders. S&F's main metatheoretical thrust is that in view of certain complex neurological findings (in conjunction with certain relatively simple behavioural manipulations) the computational metaphor, although not to be scrapped outright, needs a serious overhaul. An alternative view is that the computational metaphor, although not immune to attack, cannot seriously be threatened by this *kind* of attack, for the following reasons:

Clarity and chaos. It is difficult to appreciate S&F's argument clearly because of a marked tendency to switch chaotically between different levels of discourse. On the one hand they appear to support naïve materialism at the *neural* level; the "theory can be tested . . . by physiological experiments; it is not *merely* computational" (Sect. 3.6, para. 5; italics added). Of course, computational theorists are not constrained by the merely physiological, but that is beside the point. On the other

hand we are told that we will not understand "large-scale neural circuitry . . . until the mathematical theorists . . . have caught up with experimentalists, or until engineers have built hardware models . . . and determined whether they behave the way parts of brains do" (Sect. 3.6, para. 6). What does this imply for the bedrock status of the neural/physiological? Are hardware models somehow more convincing simulations than "mere" computer simulations (hardware *and* software)? Perhaps, because there is a subsequent favourable reference to the simple mechanical models of the early cyberneticists. And if we change a few substantives in the quotation, it sounds suspiciously like what computationists are doing for cognition anyway – that is, *they* will never understand complex cognitive processes until appropriate formal languages have been developed for their description and they have been successfully simulated on computers. Finally, what is one to make of "the *brain* should be viewed as a self-organized *process*" (Sect. 4.1, para. 4; italics added); this is symptomatic of the general confusion over structure and function, description and explanation, computers and computation.

Metaphor and simile. Part of this confusion arises from the assumption that the *computational* metaphor is just the *computer* metaphor. S&F's characterisation of the former is essentially a description of a rather basic computer; I doubt whether many computational theorists would wish to defend such a description as adequately capturing the features of a complex organism. The computational metaphor is usually seen as much more abstract; indeed, it has often been said that, since any process can be construed as a computational process, the metaphor is tainted with overgenerality, tautology, irrefutability, or emptiness. [See Pylyshyn: "Computation and Cognition" *BBS* 3(1) 1980.]

The main problem with metaphorical assertions is that they can be taken as literally false (it doesn't rain "cats and dogs"). Why, then, can't they be taken as literally true, as saying something about "reality" ("information is held in short-term memory and transferred to a long-term store")? Now, whereas metaphor misleadingly implies a kind of identity, simile makes the weaker implication of resemblance. If the rain is "*like* lead shot" we can at least ask, "In what respect?" There has been much needless debate because relations of simple resemblance have been stated or understood as something stronger. How can a model of discrimination be explanatory if it contains a primitive element that discriminates? If we say that a suitably programmed computer "understands questions," does it *really* understand? How can we have visual images when there cannot *really* be pictures in our heads? There are many such examples. On the other hand, when simile is seen for what it is, such problems do not arise; hydraulic ethological models were never seriously criticised because there seem to be no pipes and valves in the nervous system, and Freudian hypotheses are not falsified by pointing to the absence of three interacting figurines in the skull. [See also Hoyle: "The Scope of Neuroethology" *BBS* 7(3) 1984.]

The computational metaphor has strength, flexibility, and appeal because it is not really a metaphor, it is just a simile; and if the functional resemblance between two systems is sufficiently convincing, that is all that matters. But S&F are not convinced. Why? They look into the nervous system (albeit very indirectly) and find no symbols, only dynamic neural patterns. Is this surprising? Symbols are in the eye of the beholder. One can look at this page and find no symbols, only patterns of grey. The popular dogma that brains and computers deal in symbols is misleading; their currency is electricity. By saying that a device manipulates symbols we are attributing intentionality to it – how else would it know what things were or were not "symbols"? But if dynamic neural patterns are to be *discriminated*, then surely they can be named, formalised, *computed*? Otherwise we would only see chaos, in its everyday sense.

It would be foolish to suggest that neural dynamics or statics are completely irrelevant for an understanding of mind and

behaviour, and equally foolish to suggest that the computational metaphor is impregnable. If S&F are simply claiming that their data require different kinds of computations, then this is unexceptionable, but they seem to be simultaneously attacking a straw man and trying to throw the baby out with the bath water (not literally, of course). Principled attacks on the computational approach are likely to be top-down in terms of intentional and experiential arguments (Dreyfus 1972; Caudl & Shotter 1977; Searle 1980). And since computational theorists handle intentionality at worst trivially and at best controversially, I fail to see how neuroscientists could even begin to tackle the issues.

When the "chaos" is too chaotic and the "limit cycles" too limited, the mind boggles and the brain (model) flounders

Michael A. Corner and Andre J. Noest

Netherlands Institute for Brain Research, 1105 AZ Amsterdam, The Netherlands

Let's start from the beginning. To begin with, the olfactory bulb must respond to each odor the mucosa is capable of discriminating with some sort of specific pattern of excitation, presumably derived from the ontogenetically determined distribution of osmochemical receptor specificities. If each class of receptor cells is spread out over a large enough area of the olfactory mucosa, a widespread projection of odor-specific sensory volleys can be guaranteed even if, as in other sensory systems, nearest neighbor relations are largely preserved in the central projections. In this sense Skarda & Freeman's (S&F's) tantalizingly brief statement in Sect. 3.6 about the absence of a requirement for ontogenetic constraints on afferent topography is true enough, as far as it goes (provided, of course, that each receptor's terminal field does not encroach too much upon the territory of its neighbors). The next step in olfactory discrimination would be the evocation of a diffuse polynuclear oscillation in the bulb during the inhalation phase of each sniff in a motivated animal. This "something is out there" carrier waveform is manifested as a "chaotic" broad-band EEG signal within the gamma range of frequencies (ca. 40–80 cps). The dual effect of a smell – to provide sensory information together with a nonspecific signal preparing the brain for dealing with it – thus resembles, in general terms, the classical picture of a "reticular arousal system" linked to sensory projections to the neocortex. In the latter case, however, rather than high-frequency waves being triggered, low-frequency waves (EEG alpha and delta bands) become suppressed during arousal. Alerting responses in the septohippocampal system in turn consist of a synchronized neuronal oscillation, but one which is much slower (in the EEG theta band) than the one found in the olfactory system. Why these differences?

S&F's suggestive attempt to generalize their paleocortical model by postulating the existence of similar waves (i.e., in the EEG "gamma" range) in the neocortex (that have gone undetected owing to cytoarchitectonic differences between the two structures) fails to reach our plausibility threshold. The laminar organization of neocortical tissues would appear to be eminently suitable, despite its relative complexity, for detecting even weakly synchronized fluctuations of neuronal activity. An important task facing any theory which aspires, albeit implicitly, to providing a general explanation of the biological significance of "brain waves" must surely be to explain the appearance of prominent cortical oscillations – alpha rhythms, "spindling," delta waves – precisely at those times (ranging from drowsiness to deep sleep) when sensory processes appear to be at a minimum even with respect to internal sources – as in dreaming. On the other hand, the basic notion of widely synchronized neuronal carrier waves that become "destabilized" by afferent input

(see Sect. 4.5), in a spatially distinctive manner for each discriminable stimulus, is by no means excluded by differences among brain regions displaying the precise characteristics of these (chaotic) waves. Perhaps the major challenge for S&F's model, therefore, will be to account for the olfactory system displaying the very EEG waveform and amplitude-pattern that have actually been observed. By the same token, light needs to be shed on the possible significance of the low-amplitude, highly chaotic, "background" EEG present between sniffs (as well as continuously in a nonmotivated animal). This is the activity, after all, which the authors believe (see Sect. 3.2) constitutes "the elemental phenomenon . . . in all of brain physiology"(!). Disappointingly, it is almost totally neglected thereafter, although this omission became apparent to us only after our realization that the broad-band EEG gamma waves seen during (motivated) inhalation in a naïve animal were *not*, in fact, what was meant by the term "background" activity as used in the target article!

Odor-specific differences are reported to become overtly manifest in the multichannel olfactory EEG only after proper reinforcement has taken place. Spatially distinctive amplitude patterns are then detectable, taking the form of extensive limit-cycle activity in the bulb, contained within a relatively narrow band of EEG gamma frequencies. These patterns presumably reflect the magnification of preexisting differences in the spatial distribution of afferent signals and evoked synaptic activities in the bulb, without which no distinction among various inhaled (unconditioned) odors could have been made in the first place. What, then, needs to be "learned" about such signals or (as S&F would put it) to be added to the animal's smell repertoire? Nothing else, surely, than that the odor in question has acquired a particular behavioral significance: *eat* it, *jump* it, *avoid* it, and so on. This being the case, isn't it possible – even likely – that each recognizable new EEG pattern carries information not about the input but, rather, about the *output* side of the olfactory loop (i.e., the motor response system to which the stimulus has become linked by virtue of conditioning)? We're very much interested in knowing, therefore, exactly how many of these distinctive spatial patterns have in fact been identified, and whether two odors with more or less the same "meaning" for the organism would stand much chance of being discriminated on the basis of EEG analysis.

Finally, serious semantic ambiguities have arisen in the course of our attempt to understand the more strictly mathematical aspects of S&F's paper. After satisfactorily dispensing with the straw man of digital computers as useful for modeling any kind of brain, the authors proceed to find fault with "connectionist" models because of their current shortcomings in the light of recent neurophysiological findings. But in what sense is the Freeman model – not the lumped (i.e., spatially averaged) version described here, which, by definition, is incapable of even beginning to deal with the spatial EEG patterns on which the whole theory rests, but the promised but still preliminary distributed model – itself not a connectionist model? In the absence of any definition of a qualitatively new class of models incorporating features that are inherently absent in a connectionist approach, S&F's scheme must be considered as constituting simply a possible improvement *within* that category. If we then try to pinpoint what their precise suggestions for the incorporation of new features are, we are unable to find any satisfactory starting point for carrying out the proposed improvements. Several of the deficiencies attributed to existing models, such as failure to incorporate inhibition, asymmetric synapses, or endogenous noise in the system, fail to do justice to the state of the art in this field. Even if the next step were to entail the introduction of coupled limit-cycle oscillators, distributed models involving sheets of interacting circuits (each resembling the basic one in the lumped model presented by Freeman & Skarda) have in fact already been studied extensively.¹

Even in structurally homogeneous variants of such models, oscillating activity can (among many other possibilities) become ordered in spatially inhomogeneous, nonperiodic patterns. These can usually be characterized by topologically conserved phase patterns involving "vortex" – or "string" – singularities embedded in a smooth phase-field. It can be predicted that in case of spatial smoothing over a scale larger than the size of the vortices, such *phase* patterns would appear instead as spatially nonuniform *amplitude* patterns associated with a smooth phase-field. If the present data turn out not to be explicable along these lines, then it seems logical to assume some form of "pinning" of the oscillatory patterns by structural disorder. It is plausible to suppose that each of the many possible distinct patterns could then be "nucleated" by the appropriate set of incoming stimuli. Developing such conjectures into testable theories will probably require that investigators start delving into the complexities of structurally inhomogeneous models. In view of the many possible ways of generalizing the existing ones, it would be extremely helpful if experimentalists attempted to specify, as precisely as possible, the lessons to be learned (for example, from olfactory cortex physiology) that would allow such improved models to be developed.

It is wonderful for psycho(physis)logists to master the mathematics of cooperative networks, and to try to apply this knowledge to the unraveling of the deepest (or even the superficial) mysteries of the brain, but the required conceptual underpinnings for such flights into higher spheres must not be neglected. In our opinion, much more attention needs to be devoted to such fundamental things as clarity of definitions, explication of assumptions, rigor in logical structure, and completeness in the consideration of relevant theoretical and empirical material.

NOTE

1. There exists a considerable body of literature on spatially distributed, coupled limit-cycle oscillations. Good lists of core references are cited by Oono and Kohmoto (1985) and Winfree (1980).

On the differences between cognitive and noncognitive systems

D. C. Earle

Department of Psychology, Washington Singer Laboratories, University of Exeter, Exeter EX4 4QG, England

Skarda & Freeman (S&F) interpret their findings as supporting the proposal that brain function is a self-organized process of adaptive interaction with the environment, a process to be conceived in connectionist terms and involving parallel distributed processing. These views are set in opposition to the proposal that the brain is a rule-driven and controlled system solving problems, completing patterns, and forming hypotheses by manipulating symbols.

Two separate issues are conflated here. The first concerns the question whether the appropriate model for the brain is the connectionist model, with distributed parallel processing, or that taken from the digital computer, with a limited-capacity central processor and a sequential organization. This question is separate from whether the brain is a symbol-manipulating and rule-driven problem-solving device. The former is a question about the functional architecture of the brain, whereas the latter is the question whether or not the brain is an information-processing device. A distributed information-processing system may implement rules, complete patterns, manipulate symbols, and, if need be, formulate hypotheses. Consider, for example, a distributed information-processing system that takes as its input a symbolic representation, performs a transformation, and then outputs a different symbolic representation. Such a system is exemplified by certain implementations of the cooperative stereo-matching algorithm proposed by

Marr and Poggio (1976). An information-processing system may be described at the highest level in terms of its computational theory; however, the computational theory may be realised in devices using different functional architectures – that is, devices with distributed parallel processing or sequential processing.

To adopt the proposal that the brain is an information-processing and symbol-manipulating system is a methodological decision of its proponents, and as such constitutes adherence to a particular research programme. An information process may be realised in a neural system or a computer, whether it is sequentially organized or performs distributed parallel processing, but on the higher level of analysis it remains an information-processing system. As such, it should be described using informational terms – that is, rules, symbols, representations, and the language of information-processing operations (e.g., detection and discrimination). Information processes are not to be conceived directly in neurophysiological terms or in the terminology of electronics. If, for example, one says of a certain cell in the visual cortex that it is a bar detector or that it makes a measurement on the image, then one describes that cell in information-processing terms (insofar as detection and measurement are information-processing operations). At this level of description the output of such a cell and its interactions with any neighbouring cells are symbolic in that they represent a detection or its absence, or a measurement.

If the information-processing paradigm is not adopted, then the appropriate terminology is not that of symbols, rules, and so on, but a description in terms of whatever are now judged to be the intrinsic properties of the system being described. In the case of a neural system, these may be synaptic connections, inhibition, fatigue, electrical impulses, and perhaps chaos, attractors, and repellers. Thus, an account of the visual tilt aftereffect can be given in terms of differential fatigue of cortical cells with orientated receptive fields without recourse to the language of information-processing systems; it can be given in terms of the intrinsic properties of the neural substrate. In such a case there may be no basis for adopting the information-processing paradigm and giving an account using the notion of bar detectors – although, in principle, the account could be phrased in these terms were that paradigm to be adopted. It may not be necessary to use the language and concepts of information processing to give an account of the tilt aftereffect, but when considering the correspondence problem in stereopsis (Marr 1982) there may be considerable advantages in using such a language. The connectionist movement offers two fundamental challenges: First, a different functional architecture for information-processing systems is proposed. Second, and separately, the connectionists claim to provide a way of describing the behavior of aggregates of processors without the assumption that the processing is an informational one.

The problem of distinguishing between cognitive and non-cognitive self-organising and distributed processing systems may now be viewed differently from the position adopted by S&F. I propose that the critical property distinguishing cognitive from noncognitive systems is not adaptivity, but information processing. One would not want to say of a weather system that it is an information-processing system, and one would explain its behaviour in terms of the intrinsic properties of the system – that is, the pressure and temperature of air masses, humidity levels, turbulence, and so on. A variety of homeostatic and adaptive devices (e.g., thermostats and eucalyptus trees) can also be described in terms of their intrinsic properties without appeal to informational concepts. Perhaps the major challenge of the connectionist movement in relation to psychology is that, although not necessitating a noncognitivist stance, it nevertheless promises to provide a noncognitive account of complex behaviour.

Finally, it is to be noted that S&F have given a connectionist account only of neural activity in the olfactory bulb. Their claim,

however, is to have given such an account of odour recognition and discrimination, and this is a different matter. As they are at pains to emphasize, the connectionist processing that they describe for the olfactory bulb must be linked to the motor system to enable interaction with the environment – that is, discriminative behaviour. To this end, a particular pattern of neural activity in the bulb must serve as the condition for a condition-action link. One difficulty here is that a distributed parallel processing module may be embedded in a more complex hybrid and controlled system with a sequential organization. Furthermore, a condition-action link can be interpreted as a rule in an information system or as a direct neural pathway in a noncognitive system. Skarda & Freeman attempt to draw conclusions concerning the brain as a whole on the basis of the study of only a small part of the neural substrate of a discriminative interaction with the environment.

The virtues of chaos

Alan Garfinkel

Department of Kinesiology, University of California at Los Angeles, Los Angeles, Calif. 90024

Only recently (Lorenz 1963) was it realized that deterministic systems can display behavior that appears random. This phenomenon, called “chaos,” offers a new approach to modeling erratic processes. It should be stressed that the chaos that arises in deterministic systems is not total chaos, but rather is controlled and bounded, and has definite qualitative form. It also differs from ordinary random behavior in that it is low-dimensional, whereas traditional “noise” arises from the central limit theorem, which predicts a normal distribution from the addition of a large number of independent contributions.

Skarda & Freeman (S&F) propose that the background EEG in the olfactory bulb is chaotic. The principal evidence for this claim is their report of calculations of the “dimension” as lying between 4 and 7. Such calculations of apparent dimension are one way of distinguishing chaos from noise, although there are difficulties and pitfalls in this approach (see especially Grassberger, 1986, for a discussion of fallacious calculations).

But the calculation of dimension is only one way of distinguishing chaos from noise, and it suffers from being just a number. Methods like attractor reconstruction and Poincaré sections (Froehling, Crutchfield, Farmer, Packard & Shaw 1981) have the additional advantage that they give qualitative pictures of the behavior and of the form of its underlying mechanisms. Such information is much deeper; the dictum here is that “quantitative is just poor qualitative.” See Roux, McCormick, and Swinney (1981) and Farmer, Hart, and Weidman (1982) for applications of these methods to chemical chaos and fluid turbulence.

Once one has established the fact that a given phenomenon is a chaotic process, the next question is: What is chaos doing there? S&F suggest that it is playing a functional role, an idea that is something of an about-face for chaos. Most writers on the subject tend to assume that chaos is something bad; the proposed examples of chaos in physiology, such as cardiac and respiratory arrhythmias (Glass & Mackey 1979), would support this view. But chaotic behavior can also be functional and adaptive. Consider the chaos of fluid turbulence. Turbulent fluids have useful properties that are not found in nonturbulent states. For example, the mixing properties of turbulence greatly increase the fluid’s ability to take imposed heat or movement and equipartition it out.

A striking example of the functional role of chaos can be found in a study of population dynamics by Auslander, Cuchenheimer, and Oster (1978). They study a model of a coevolving host/parasite system and find that, for certain ranges of key

parameters, the host species displays chaotic variations in population level over time. They interpret this by suggesting that "for a host population being pursued by a coevolving predator, it is surely adaptive to maintain a demographic and genetic pattern as 'untrackable' to the parasite as possible" (p. 290).

It may well be that in physiology, chaotic behavior can be quite useful, serving to randomize a system in cases where regular behavior would be damaging. A case in point might be the normal cortical EEG. It is interesting that many writers have suggested that epileptic seizures might be examples of chaos. In fact, the opposite seems to be true: In seizures, the EEG becomes regular and periodic, and it is the normal ("desynchronized") EEG that is irregular. Given the undesirability of periodic cortical behavior, it is reasonable to suppose that the nervous system has evolved a reliable mechanism to desynchronize the EEG. As an example of the utility of such "active desynchronization," consider the behavior of a platoon of soldiers crossing a bridge. Since periodic behavior (marching in ranks) might set the bridge into destructive resonant oscillation, the soldiers "break ranks." It is likely that the nervous system can effect a similar active desynchronization, in situations where randomness is too important to be left to chance. In the case of the soldiers, there is a commanding officer who gives the order to desynchronize. Here the difference between high-dimensional noise and low-dimensional chaos becomes crucial. 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: Only a few parameters need be altered to move the system into and out of chaos.

A similar phenomenon may occur in muscle activation. If individual motor units were to fire periodically, they might tend to synchronize, producing undesirable tremor. Hence, there may well be an active desynchronization mechanism in sustained contraction that "spreads out" the motor unit timings to fill the time interval of the activity.

In general, it may be that for all oscillatory processes in physiology, a perfectly periodic oscillation is undesirable. Chaos could here play the role of introducing a useful wobble into the period or amplitude, while retaining the overall form of the process.

Stable self-organization of sensory recognition codes: Is chaos necessary?

Stephen Grossberg

Center for Adaptive Systems, Boston University, Boston, Mass. 02215

Freeman and his colleagues have developed one of the classical experimental and modeling paradigms of neurobiology through a remarkable synthesis of technical virtuosity, physical intuition, and intellectual courage. Their systematic approach has led them to articulate a number of fundamental problems concerning the self-organization of sensory codes and to propose possible approaches to the solutions of these problems. Due to these characteristics, data and modeling ideas from the Freeman school provide one of the best vertebrate sources of quantitative results about interactions between cortical sensory representations and their appetitive modulation, and have therefore played a valuable role in testing the principles and mechanisms of adaptive resonance theory (Grossberg 1982a).

Even in such a systematically explored neural paradigm, definitive data are more the exception than the rule. In the absence of data that afford a unique specification of generative neural mechanisms, a number of theoretical tools can be invoked to impose additional constraints. Two such tools are mathematical and simulation analyses of the emergent properties of those model neural systems that are consistent with basic neural organizational principles and the data at hand.

Such results are briefly reviewed here to help weigh the key hypotheses of Skarda & Freeman (S&F).

After reviewing fundamental facts about spatial pattern coding by temporally entrained waveforms and the role of associative learning and expectancies, S&F focus on their central "conjecture that chaotic activity provides a way of exercising neurons that is guaranteed not to lead to cyclic entrainment or to spatially structured activity. It also allows rapid and unbiased access to every limit cycle attractor on every inhalation, so that the entire repertoire of learned discriminanda is available to the animal at all times for instantaneous access" (Sect. 3.4, para. 3). They propose this hypothesis in an unusually strong form, going on to claim that "without chaotic behavior the neural system cannot add a new odor to its repertoire of learned odors" (Sect. 4.4, para. 1).

S&F's hypothesis raises the difficult issue that a data phenomenon, despite its *correlation* with a particular functional property, may not be *necessary* to achieve that functional property. When this is true, it is not possible to assert that the system has been *designed* to generate the property for that functional purpose. One can defeat the claim that the property in question is necessary by providing a mathematical counterexample of its necessity.

Gail Carpenter and I (Carpenter & Grossberg 1987) have done just that. We have completely analyzed an explicit example of an adaptive resonance theory architecture, called ART 1, which shares many features with those of the Freeman data, and we have mathematically proved that this architecture has the following properties. ART 1 can self-organize, self-stabilize, and self-scale a sensory recognition code in response to an arbitrary, possibly infinite, list of binary input patterns. During code learning, it carries out an efficient self-adjusting memory search; after learning self-stabilizes, recognition occurs without search by direct access to the optimal learned recognition code. The course of learning is, moreover, remarkably stable; all adaptive weights, or long-term memory traces, oscillate at most once through time due to their dynamic buffering by system interactions.

Such ART architectures have been explicitly designed to provide "the system with continued open-endedness and readiness to respond to completely novel as well as to familiar input, without the requirement for an exhaustive memory search" (S&F, Sect. 4.4., para. 2). To discover systems capable of coping with this "stability-plasticity dilemma," all the operations of ART systems have been developed to explain and predict difficult parametric data from a number of behavioral and neural paradigms: for example, Cohen and Grossberg (1986), Grossberg (1987a,b), Grossberg and Stone (1986).

The key point for present purposes is that chaos plays no role in the extremely flexible and powerful learning and recognition performance of such a system. Hence chaos is not *necessary* to achieve the type of competence that has been uniquely ascribed to it by S&F.

This argument does not deny that chaos has been measured in S&F's experiments. In fact, one of the important mathematical issues in neural network theory concerns the manner in which parameter changes within a single neural model can cause bifurcations between point attractors, limit cycle attractors, and bursting or chaotic attractors. In a book to which Freeman and I both contributed, I reviewed results concerning how suitable parameter changes could cause a cooperative-competitive feedback network to bifurcate from one with point attractors into one with limit cycle attractors of a type (standing waves) that could support an olfactory code (Grossberg 1981). It was there suggested how a parameter called the quenching threshold (QT) could modulate the olfactory bulb's excitability in phase with the breath cycle; and it had been known for some time (Elliass & Grossberg 1975) how such gain changes could cause a Hopf bifurcation from a point attractor into a limit cycle. The bifurcation described by S&F from a low-energy state to a high-energy

state with a narrow temporal spectral distribution is clarified by such results. So too is their observation that during early inhalation, intrinsic interaction strength is low, since one way to alter the QT is by altering the gain of system interactions, as occurs in ART 1 through its *attentional gain control* channel (Carpenter & Grossberg, 1987). When attentional gain control is low, cell populations can become decoupled.

These relationships between point attractors and limit cycle attractors delineate a family of models, all of which can support similar functional coding properties, with or without chaos. A formal model with such functional properties can also possess a tonically active point equilibrium or chaotic attractor in its rest state. Such a state of tonic activation can support one or more basic functional properties [e.g., maintaining a baseline activity that can be excited or inhibited without a loss of sensitivity, feeding signals into habituating chemical transmitters that can compensate for spatial fluctuations in the basal activation level and therefore keep the tissue in a spatially unbiased state (Grossberg 1983), or driving antagonistic rebounds in response to sudden offsets of sensory inputs (Grossberg 1980)]. All these functions can be carried out equally well by point or chaotic attractors. On the other hand, in a state of tonic activity but low attentional gain control, a physically realized network of cells can exhibit small but complex, even chaotic, fluctuations.

In summary, just as one can conceive of slime mold aggregations that proceed continuously or in a pulsatile fashion through time as parametric variations of a single model, so too can one envisage a sensory coding model in which point or chaotic attractors support similar functional characteristics. Thus, although the issues raised by S&F are important ones for understanding cortical design, further argument is needed to support their strong claim for the *necessity* of chaos to achieve key functional coding properties.

S&F have also raised the legitimate challenge that "no claim for firm and substantial understanding of large-scale neural circuitry can be advanced until the mathematical theorists of distributed dissipative systems have caught up with experimentalists, or until engineers have built hardware models" (Sect. 3.6, para. 6). The Carpenter-Grossberg (1987) theorems have, in fact, provided such mathematical guarantees about the ART 1 architecture, and these guarantees have encouraged engineers to start building an ART 1 chip in hardware. The kind of functional competence Skarda & Freeman have seen in their data is thus already helping to define the technological products of a biologically derived artificial intelligence.

ACKNOWLEDGMENT

Supported in part by the Air Force Office of Scientific Research (AFOSR 85-0149) and the National Science Foundation (NSF IRI-84-17756).

Is chaos the only alternative to rigidity?

Daniel S. Levine

Department of Mathematics, University of Texas at Arlington, Arlington, Tex. 76019

The Freeman laboratory continues to do exciting work in the neural representations of olfactory stimuli, as it has for almost 20 years. As always, the Skarda & Freeman (S&F) work is pioneering, both technically and philosophically. The changes in EEG patterns from a neutral to a reinforcing odorant stimulus are particularly significant. From the modeler's perspective, however, I retain some skepticism about the philosophical role of chaos that S&F propound.

S&F write that "the process of state change leading to the unstructured chaotic domain is essential for preventing convergence to previously learned patterns, and hence for the emergence of new patterned activity" (Sect. 4.4, para. 2). Although the authors state this most strongly for the olfactory

bulb, they clearly hope to apply the same principle to other sensory modalities, and have adduced some evidence for it in the visual system.

In neural models of the "connectionist" variety, based on nonlinear differential equations, the capacity to respond to both novel and familiar inputs can exist even in the absence of chaos. In fact, learning often transforms novel inputs to familiar ones, with the consequent change in response properties. A series of articles by Grossberg (in particular, 1975; 1980; 1982b) discusses a striving for balance between an attentional system (which biases the network's responses toward previously learned inputs) and an arousal system (which enables the network to overcome the attention system's rigidity when important new events occur). One mechanism for responding to novel events in these networks is the activity of populations of "mismatch detectors," which are actively inhibited by *correspondence* between the activity patterns in two separate on-center off-surround fields of cell populations (such as one field representing an actual "bottom-up" stimulus event and another representing a "top-down" expectation of an event). This correspondence causes the total energy from the summation of the two fields to be sufficiently large to shut off the mismatch detector activity, which plays a role analogous to S&F's chaos. Like the chaotic EEG pattern, mismatch activity ensues if an unfamiliar pattern occurs, because the mismatch is not inhibited by correspondence. Hence, in Grossberg's model as in S&F's work, "an animal can classify an odorant as 'novel' with no greater delay than for the classification of any known sample" (S&F, Sect. 3.4, para. 3).

There is obviously no complete isomorphism between model network "minimal anatomies" (Grossberg 1975) and real neuroanatomies. S&F's point about the limitations of the connectionist model of Hopfield and Tank (1986) is well taken (Sect. 4.4, para. 4). Both in that model and in the more general model of Cohen and Grossberg (1983), theorems show that the network always converges to an equilibrium state representing a "decision" about short-term pattern storage. As S&F rightly point out, such behavior is too circumscribed for the actual nervous system. In fact, in real-time simulations of behavioral data, networks of that variety typically model only the short-time dynamics of a part of an entire network that includes both associative and competitive parts. An example of such real-time simulation occurs in the Pavlovian conditioning model of Grossberg and Levine (submitted) [also summarized by Levine (1986)].

Hence there is room for many years of research on how established neural network theories of specific processes concatenate into large-scale systems that actually reproduce significant data, in the olfactory cortex and elsewhere. Obviously, the chaotic EEG patterns in the resting olfactory bulb are important to the theory of that area. Whether the chaos is essential to the purpose of the bulb's function or epiphenomenal to other things that are essential, I cannot hazard a guess. The answer will be related to that of the larger unanswered question concerning what the EEG measures in general!

Chaos in brains: Fad or insight?

Donald H. Perkel

Theoretical Neurobiology Facility, Department of Psychobiology and Department of Physiology and Biophysics, University of California, Irvine, Calif. 92717

The brain sciences, in their more reflective phases, are notorious for their immersion in analogy and metaphor (e.g., Arbib 1972). Traditional brain metaphors arise from technology. The nineteenth-century analogy between neuronal processes and undersea cables survives as modern cable theory. Subsequent

technological metaphors have been based on telegraphic networks, telephone exchanges, control systems (Ashby 1952), digital computers (von Neumann 1958), holograms, and nonlinear networks (Hopfield & Tank 1986; Rosenblatt 1962; Rumelhart, Hinton & Williams 1986). Each of these metaphors has contributed valuable insights, some more than others; none provides a global theory of brain function.

Mathematical structures have also served as neural metaphors. Probabilistic examples include random-walk models for impulse-interval distributions (Fienberg 1974; Gerstein & Mandelbrot 1964; Sampath & Srinivasan 1977), stochastic point-process models of nerve-impulse sequences (Moore, Perkel & Segundo 1966; Perkel, Gerstein & Moore 1967a; 1967b), and the binomial model for quantal release of neurotransmitter (del Castillo & Katz 1954; Zucker 1973). Other primarily mathematical theories include the formal neuron model of McCulloch and Pitts (1943), interacting oscillator theories of the EEG, thermodynamically inspired theories of interacting populations of nerve cells (Cowan 1968), information theory as a paradigm for brain function, tensors as the basis of cerebellar function (Pellionisz & Llinás 1979), and the "trion" theory of cortical cell assemblies (Shaw, Silverman & Pearson 1985), essentially a probabilistic cellular automaton (Wolfram 1984).

Not all of these mathematical metaphors have fared well in the neuroscientific community. Random-walk models for impulse-interval distributions make nonunique predictions. The strict binomial model for neurotransmitter release yields misleading interpretations of experimental data (Brown, Perkel & Feldman 1976). Other mathematical models have been criticized on the grounds that the mathematical structure has dictated the biological assumptions or that the theory was leading the data.

Recently, much attention has been paid to the modern treatment of nonlinear differential equations, including catastrophe theory, bifurcation theory, Poincaré maps, strange attractors, "chaos," and fractals. Biological applications have abounded, sparked by May's (1976) demonstration of chaotic behavior in population dynamics. Bifurcation theory has been applied to excitable cells (Chay & Rinzel 1985). Skarda & Freeman (S&F) make broad claims about the explanatory role of bifurcations and the emergence of "chaos" in the functioning of the olfactory bulb. Similar claims have been advanced for activity in invertebrate ganglia (Mpitsof & Cohan 1986) and in cardiac arrhythmias (Mandell 1986), among others.

The question that immediately arises is whether the biological phenomena themselves dictate or justify the theory's mathematical structures. The alternative is that the beauty, versatility, and power of the mathematical approach may have led its aficionado to find areas of application in the spirit of the proverbial small boy with a hammer, who discovers an entire world in need of pounding. Is bifurcation theory merely a trendy framework for a Procrustean approach to nervous-system function? Does it make any more sense to say that the olfactory bulb makes chaos to make sense of the world of smell than it does to say that the cerebellum is a tensor, or that the hippocampus is a map, or that the visual system is a Fourier transformer, or that cognitive processes are executions of computer programs? Is the theory of familiar and strange attractors a natural way of looking at neurobiological phenomena – at the olfactory bulb in particular – or is it a method in search of a roosting place?

At the cellular level, the use of bifurcation theory by Chay and Rinzel (1985) clarifies the behavior of their system in a plausible and rewarding way; it enriches our insight. However, the bulb is immeasurably more complex, far less perfectly characterized, and harder to measure than the single cell; bifurcation analysis of the bulb is necessarily more risky, less readily quantifiable, and more subject to distortion.

Assuming that surface EEG measurements sufficiently well represent mitral-cell firing rates, what S&F have sketched is not a theory of odor recognition and learning, or of olfactory bulb

function, but rather an outline of a research program to produce and refine such theories. Their experimental findings, although far from conclusive, in fact make their argument plausible, in the context of the behavior of other nonlinear dynamic systems.

S&F correctly point out that connectionist models can generate chaotic behavior if artificial constraints on connectivity are lifted. A serious problem, however, remains: How does the system read out the information – that is, the identity of a familiar odorant – when its "representation" is so dynamic and volatile? The answer must lie in the anatomy and physiology of the bulb and more central structures, but the working principles of specific odorant identification remain to be elucidated.

Do the operating principles of the olfactory system hold for other sensory systems that have highly topographic anatomical representations? It may be that widespread chaos and self-organization are peculiar to the olfactory system or the brain stem, and that topographic systems "use" chaos in a much more restricted fashion.

Inhibition, as S&F point out, is essential to the operation of the system. Unaccountably, they mention the strengthening of excitatory synapses but not inhibitory synapses, although Wilson, Sullivan, and Leon (1985) describe increased inhibition in mitral cells after olfactory learning. It seems prudent to impute plasticity to inhibitory synapses as well.

S&F lament the weakness of the purely mathematical methods. The inescapable remedy is to mount a series of increasingly realistic, large-scale simulations of the system. The chief contribution of digital computers to theoretical neurobiology may be as tools for analysis and synthesis, rather than as marginally appropriate metaphors.

Finally, what is most attractive about S&F's theoretical approach is the biological flavor of its predictions. The picture of a spontaneously active bulb, goaded by sensory input into chaotic-appearing nonrecurring spatiotemporal patterns of activity, was sketched almost half a century ago: "millions of flashing shuttles weave a dissolving pattern, though never an abiding one; a shifting harmony of subpatterns" – the "enchanted loom" of Sherrington (1940; rev. ed. 1953, p. 178). When the skeletal theory has been fleshed out with more fine-grained experimental evidence and correspondingly realistic simulation studies, it may well be that bifurcation theory and chaos, arising out of "connectionist" models, may provide a cohesive, unifying, and apt theory for widespread aspects of brain functioning.

Connectionist models as neural abstractions

Ronald Rosenfeld, David S. Touretzky, and the Boltzmann Group

Computer Science Department, Carnegie Mellon University, Pittsburgh, Pa. 15213

Skarda & Freeman's (S&F's) findings and interpretations provide strong support for the connectionist paradigm. They clearly illustrate the importance of distributed representations and dynamic system theory for understanding computation in the brain. The paper concludes by criticizing various aspects of current connectionist models. It is this criticism that we wish to address.

Connectionist models are chiefly concerned with computational aspects of cognitive phenomena. At the current stage of this research, simplicity is often preferable to biological fidelity. We realize that the brain is likely to employ mechanisms beyond our present computational taxonomy, let alone our understanding or mathematical tools, but we nonetheless believe that current models, crude though they may be, advance the understanding of cognitive systems and contribute to the emergence of a new taxonomy. One should not confuse claims about the accuracy of certain connectionist models vis-à-vis real nervous

systems with claims about their computational adequacy or scientific utility. S&F appear to have made this mistake.

S&F's target article repeatedly emphasizes the superiority of dynamic attractors over static ones, holding that connectionist models are inadequate since they do not have the former. But this is not so; a Boltzmann machine (Ackley et al. 1985) annealed down to a temperature slightly above its freezing point is manifesting a dynamic attractor state very similar to the one advocated by S&F. More important, the target article fails to demonstrate any computational advantage of dynamic models. Connectionist models are abstractions. Stationary patterns of activity in these models need not correspond to stationary patterns in the brain, just as connectionist units and their weighted connections need not correspond one-for-one with real neurons and synapses. Connectionists are perfectly happy to stipulate that the stable states of a Hopfield net (Hopfield 1982) or a Boltzmann machine are abstractions of dynamic attractors in the brain. We will abandon models with simple point attractors only if dynamic models can be shown to have useful computational properties that static ones lack. We have not yet seen the evidence that could support such a claim.

S&F maintain that chaotic behavior is essential for learning, but they do not make clear what role chaos is supposed to play in the learning that takes place in the rabbit olfactory bulb. The target article claims that a chaotic well – a “don't-know” state – is a prerequisite for the system to learn to recognize new odor categories. But which of the characteristics of chaos are necessary to the role it plays in generating new attractors, and which are irrelevant? S&F's article does not answer this key question.

S&F further criticize connectionist models because of their need to be externally reset after reaching a stable state. But the olfactory bulb does in fact settle into a single (albeit dynamic) state that is computationally equivalent to a corner of a hypercube; and it does not spontaneously escape from one dynamic attractor to other interesting ones. The return to the chaotic well (cf. the center of the hypercube) that takes place at exhalation in the rabbit appears to be precisely a forced reset action.

S&F next advise connectionists to give up the view of neural networks as pattern completion devices. They maintain that no pattern completion activity takes place in the olfactory bulb, since its output is a coherent global state generated from within, not merely a completed pattern within one nerve cell assembly (NCA). But to say that no pattern completion takes place in the olfactory bulb is to mix levels of description. Receptor cells send their pulses to the olfactory bulb, which in turn settles into a dynamically stable state – one of several preexisting possibilities. This is precisely what pattern completion is about! Stationary pattern completion activity in connectionist models is an abstraction. It need not correspond to stationary pattern completion in the brain. On the other hand, the “destabilization” paradigm advocated by S&F is merely a metaphor, and will remain so until it is supported by a concrete computational model.

The target article rightly points out that feedback mechanisms in the brain are far richer than those used in many connectionist models. But it also maintains that the “long delays, temporal dispersions, and spatial divergences” (Sect. 4.3, para. 2) present in the brain are *necessary* for the production of global behavior. In order to extend connectionist models to include these features, one must first have some idea of their essential role. There is no (computational) point in blindly simulating neural circuitry without first having an analytical handle on the role of the elements involved. By starting our analysis and simulation with minimal assumptions, we make sure that only essential features of the system will be admitted into our models.

Finally, we would like to point out some technical difficulties in the use of nerve cell assemblies to explain the formation of stable states. It is postulated that the NCAs are responsible for the selection of the basin to which the system bifurcates. According to this hypothesis, each NCA corresponds to a specif-

ic basin, and therefore to a specific known odor. The neurons in each NCA are supportive of one another, so that activating only some of them will cause the whole assembly to become active. How, then, is similarity between odors accounted for in this model? Do NCAs of two similar odors share neurons? If so, the presence of the first odor will activate its associated NCA. The latter will in turn activate the other NCA, irrespective of whether the odor it stands for is present. Moreover, what happens when a combination of two or more familiar odors is presented to the receptor cells? Are several NCAs activated simultaneously? What kind of basin is created, and how is it related to the basins of the component odors? What state does the system settle into eventually? The target article does not address these issues.

Chaos can be overplayed

René Thom

Institut des Hautes Études Scientifiques, 91440 Bures-sur-Yvette, France

More than a century ago the German mathematician B. Riemann, in his little-known philosophical writings, addressed the mind-body problem as follows: “When we think a given thought, then the meaning of this thought is expressed in the shape of the corresponding neurophysiological process.” It is comforting to see this old idea unearthed after hard experimental work, and put forward by Skarda & Freeman (S&F) as a major discovery. (Here, of course, “meaning” has to be understood as a nonverbal conceptualization of smells in the rabbit's psyche.) First, it seems to me, there is a gap to be filled in the findings of S&F: To what extent does the shape of the EEG amplitude on the bulb depend on the experimental procedure – in particular, on the nature of the conditioning stimulus? Would the pattern observed for a given odorant when the subject is conditioned, say, by subsequent electric shocks, be the same as the one observed when reinforcement is obtained by giving water to the thirsty subject? The rather rough model offered for the underlying general dynamics is very suggestive (S&F's Figure 11), but the idea that for each of these attractors (or rabbits' pseudoconcepts) there should exist a specific triggering NCA (nerve cell assembly) seems to me another instance of what A. N. Whitehead (1960) called the “fallacy of misplaced concreteness” (p. 11). For if, as S&F claim, there exists in principle a virtual infinity of such attractors (due to the infinite fecundity of “chaos”), then this would require an infinite number of distinct NCAs, something difficult to accept.

Here one sees clearly the limits of neurophysiological research. When one tries to describe the anatomical constraints imposed by some specific functional behavior on the physiological level, “connectionist models” ultimately mean very little – namely, that a neural mass exhibits internal symmetry of a geometric type (translation, rotation, etc.) and that this symmetry may lead to corresponding “first integrals” of the associated neural dynamics. S&F give for the word “chaos” the definition once proposed by Ruelle-Takens (1971): differential systems which display the property of sensitivity to initial data. In this they follow the present fashion, to which I do not personally subscribe. “Chaos” and “chaotic” should be reserved for systems that cannot be explicitly described either quantitatively or qualitatively (there are plenty of them). Hence, such chaotic systems have no equations. Systems defined by equations have attractors (the precise mathematical definition of which may in fact be very difficult). It is to be expected that after the present initial period of word play, people will realize that the term “chaos” has in itself very little explanatory power, as the invariants associated with the present theory – Lyapunov exponents, Hausdorff dimension, Kolmogoroff-Sinai entropy (Guckenheimer & Holmes 1983) – show little robustness in the presence of noise.

The same misuse of terminology may be seen in S&F's systematic use of "self-organizing process." By that, I suppose, they mean a process that, starting from a given set Ω of initial data, will follow a specific trajectory (Γ) to a very good approximation, at least for a given time span [or, more generally, a process exhibiting spatially invariant configurations, as for Rayleigh (1916)–Bénard (1900) convective patterns]. In such a case, the old concept of "chreod," once proposed by C. H. Waddington (1957), would do the same job, and could be given under the notion of "morphogenetic field" a very precise mathematical formulation.

All in all, I would say that the main interest of the target article lies in the physiological description of the effects of Pavlovian conditioning on a given sensory input: formation of a high-frequency peak, spatially modulated in amplitude according to a specific pattern on the bulbar surface. This dynamical finding suggests that the propagative character of Pavlovian conditioning – the "prégnance"¹ of the stimulus – could be explained as a purely dynamical effect of resonance.

NOTE

1. The French word "prégnance" was proposed by this commentator as a property of an externally perceived form that is the opposite of "saillance" (saliency).

Cognition as self-organizing process

Gerhard Werner

Department of Psychiatry, University of Pittsburgh, Pittsburgh, Pa. 15213

Cognitivists of the representation-computation persuasion could, with some justification, support their case by pointing to the absence of neurobiologically viable and conceptually consistent alternative theories. The experimental findings and the elegant interpretations presented in the target article weaken this argument substantially. Although admittedly limited to "preattentive cognition" and not incorporating aspects of attentive stimulus exploration, Skarda & Freeman's (S&F's) model contains elements of potentially more general relevance, which are awaiting further elaboration of mathematical theories of distributive, dissipative systems, and more extended validation of the correspondence between brain electrical events and stimulations according to the operational principles proposed; nor is there anywhere else in the brain evidence for the occurrence of stimulus-related high-amplitude bursts of oscillatory activity comparable to the olfactory EEG on which the interpretation of the experimental data is based. Moreover, within its own domain, the model presupposes a number of modulatory neurochemical processes and synaptic connections that await empirical confirmation before conclusive validation is possible.

Notwithstanding this current restriction in generality and conclusiveness, the concepts developed in the target article raise tantalizing issues by sketching the outlines of an internally consistent and coherent model of perception and cognition that eliminates some of the solipsistic implications of representational cognitivism.

The evidence assembled by S&F attributes a primary role to cooperative, self-organizing activity in neural structures, which can individuate situation-specific, spatiotemporal profiles of neural activity, contingent on past stimulus exposure and behavior-regulating expectancies. The conceptual implications of this position merit underscoring: History is not represented as a stored image of the past; nor is the present a mirror of the environment. Instead, environmental events are specified by states of neural activity that are the result of the neuronal system's internal organization and dynamics. In this sense, the neural structure *uses* information to *create* its own internal states, which acquire meaning: The internal states are the

neuronal system's *own symbols*, as these states stand in a regular relation to events in the world and signify potentials for action. This distinction highlights the departure from current cognitivism, for which meaning is assigned to symbols by an observer. It seems that Dretske (1986) drew a similar distinction in another context.

Once symbols are viewed as the system's own creations, any reference to representations becomes superfluous; Occam's razor can unburden us of the Trojan horse that was smuggled from the land of Artificial Intelligence into Neuroscience. Perhaps the protestations that representations exist only in the mind of the observer who jointly beholds an environment and an observed organism (brain) will at last be heard (Maturana & Varela 1980).

The overriding importance of the work reviewed by S&F lies, in my view, in the fact that it sketches the outlines of a neurologically based approach to cognition as an alternative to the tenets of current cognitivism. This in itself represents an important contribution in proposing a viable alternative to representational-computational cognitivism, and in suggesting modifications of current connectionist models. The target article sets the stage for a "pluralistic methodology," which P. Feyerabend (1975) considers a vital element in support of competitive argumentation among theories, forcing each into greater articulation, and all of them contributing to greater clarity.

Authors' Response

Physiology: Is there any other game in town?

Christine A. Skarda^a and Walter J. Freeman^b

^aCREA, Ecole Polytechnique, 75005 Paris, France and ^bDepartment of Physiology-Anatomy, University of California, Berkeley, Calif. 94720

We thank the commentators for taking the time to read, think about, and critically respond to our target article. The material we presented is diverse and difficult, despite (or perhaps in part because of) our effort to simplify it and make it accessible to researchers in other disciplines. Our exposition and our hypotheses extend from basic physiology through behavioral and cognitive theory, relying on mathematical techniques for quantitative description and prediction. The commentaries touch on all these levels and we have grouped our responses accordingly. Our overall conclusion is that our proposed view of the brain and the dynamics by which it generates behavior emerge intact from this scrutiny. However, we think that there is a problem of miscommunication that stems from failure of physiologists, psychologists, and modelers alike to follow through with careful consideration of the logical consequences of both new and long-standing findings on brain function.

Meetings, symposia, and workshops on neural networks and connectionism deriving from brain studies have now become commonplace. Yet we believe that physicists, engineers, and mathematicians have little understanding of the functional architecture of networks of real neurons, and that neural networking is just the newly derived technical capability to handle large arrays of interconnected elements with dynamic properties

that, although simple for the element, are endlessly complex for the array. Underlying this work is a weak description of a nervous system that bypasses the basic questions about the essential character and organization of brains. Physiologists and anatomists, however, are equally deficient in failing to face the epistemological and philosophical consequences of their findings and conclusions and in interpreting them in terms other than those they have inherited from reflexologists.

Our main point lies beyond level confusion, misuse of terminology, and misconstrual of the tenets of cognitive science; it is that brains don't work at all in the way everyone, including ourselves, expected them to. We asked a simple question: What is the physical form in which sensory information is registered in the olfactory bulb? The answer we found – namely, a spatial pattern of chaotic activity covering the entire olfactory bulb, involving equally all the neurons in it, and existing as a carrier wave or wave packet for a few tens of milliseconds – is orthogonal to the axes of virtually every explanatory system we are aware of. It is therefore not surprising that some of our descriptions were misunderstood, and that some of the comments should be tangential to what we wrote. The full implications take time to sink in, as do the lessons to be learned for a new technology.

What emerges from our work, as recognized most clearly by Werner, is the conclusion that the concept of "representation" (e.g., symbols, schemata, codes, maps) is unnecessary as the keystone for explaining the brain and behavior. This is because the dynamics of basins and attractors can suffice to account for behavior without recourse to mechanisms for symbol storage and invariant retrieval, "teachers" in learning devices, error-correcting feedback, comparators, correlators, associators, and the fine-grain point-to-point topographic connectionism that is required for their effectuation. The nervous system tolerates (indeed thrives on) an enormous degree of what can only be called sloppiness in its design, construction, and maintenance. This is difficult for engineers and logicians to come to terms with, even when it is dressed up as chaos; but as Garfinkel (1983), Holden (1986), Rössler (1983), Shaw (1984), and others who write about chaos have pointed out, it is a quality that makes the difference in survival between a creature with a brain in the real world and a robot that cannot function outside a controlled environment.

In sum, cognitivists have written repeatedly for some years that rule-driven, symbol-manipulating devices are the only realistic hope for mimicking and explaining behavior. We submit that the brain can do better.

Psychology: Insight or level confusion? Several commentators have raised the issue of levels of description. Have we confused different levels of description? Is our view hopelessly muddled? We think not. Barnden's comments are especially instructive in this regard.

Barnden does not contest our model of olfactory functioning as a distributed, self-organized process whose functional architecture resembles current connectionist models, but he raises two important points. First, he points out that our model of relatively low-level sensory neural mechanisms cannot be generalized without further data and argument to higher-level cognitive processes. We agree, but Barnden goes on to claim that

because "'symbol manipulationists' have in any case always presumed that low levels of perception are at least largely based on specialized mechanisms that are probably not to be regarded profitably as manipulating symbols in any conventional sense," we consequently have nothing new to offer to the debate. It is not clear which symbol manipulationists Barnden is referring to here, but we believe that he is mistaken. First, the neural processes captured by our model (after the first synapse) cannot be equated with simple transducer processes or with reflexes, both of which have traditionally been viewed as nonsymbolic. Second, it is not true that symbol manipulationists hold that input analysis, even at a relatively low level, does not involve symbolic manipulation. On the contrary, input analyses taking place after the transducer level have been considered paradigmatic examples of symbolic processing, whereas more central or higher cognitive processes have eluded analysis in similar terms. As one prominent proponent of the symbol-manipulation view says, "Input systems are computationally elaborated. Their typical function is to perform inference-like operations on representations [i.e., symbols] of impinging stimuli" [Fodor 1983, p. 83; see also multiple book review in *BBS* 8(1) 1985]. This is precisely the view that our model challenges.

Barnden also raises a more fundamental issue: Even if the lower-level sensory mechanisms could, as we claim, be explained in terms of nonsymbolic processing, what is to stop him and others from viewing the patterned output to higher cortical areas or the processes taking place in those areas in terms of symbol manipulation? Can the dynamic patterns of output from the olfactory bulb be the symbols Barnden feels compelled to look for in the brain? It is significant that Barnden is persuaded by our model to depart from the conventional view of symbols as strings of discrete bits of information that encode distinct physical properties, and to introduce what he refers to as "a certain amount of 'fuzz.'" The patterns he wants to equate with symbols are context-dependent; at best they are roughly correlated with events in the world, as he admits. As a result, crucial aspects of symbolic processing (e.g., decompositionality and inference) are jeopardized. We do have a good understanding of logical operations on conventional symbols, but we have no model for logical operations on context-dependent ones. Of course, one could still refer to these patterns as symbols, but what does the use of this term now buy us? Without the ability to perform conventional logical operations of the sort used by traditional symbolic processing, the use of the term "symbol" for the kinds of patterns we find in the olfactory system is not doing the work it did in models developed by the symbol manipulationists, and in the end it is misleading, because researchers are led to view the functional architecture of the system in a way that is not compatible with the distributed processing carried on by neural networks.

Earle grasped an important point missed by some of our commentators (Barnden and Brown): What we proposed on the basis of our data and the ensuing model is that the functional architecture of brains resembles the distributed, self-organized processes of connectionist models rather than the rule-driven symbol-manipulating processes characteristic of digital computers. Contrary to an apparently popular assumption, physiologists are not

so naïve as to conclude, because they don't find symbols floating around in the neural tissue, that the brain is not a symbol-manipulation system. We agree with Barnden and Brown that a machine-level claim would be ineffectual against this position. But our model is pitched at the level of the functional architecture of the system: Our point is that brains use the functional architecture of distributed networks similar in many ways to present-day connectionist models. Connectionist models are not plausibly conceived of as symbol-manipulating, rule-driven systems; so why, if brains use a similar form of information processing, are the latter so construed?

We disagree, however, with Earle's further claims that connectionist models promise a "noncognitive" account of behavior, and that information processing requires symbols and rules. Connectionist models are not attempts to provide a noncognitive account of behavior, if by "noncognitive" is meant not having to do with cognitive processes. Connectionist models are explicitly cognitive: Such models deal with mental processes (e.g., pattern recognition and completion, generalization, discrimination, associative memory) and the mechanisms responsible for cognition. Admittedly, our model and connectionist models in general do not appeal to rules or symbols, but this is not equivalent to the claim that they are noncognitive. As has been pointed out, connectionist models, although exhibiting regularities in processing information, do not apply rules; nor is the "currency" of such systems symbols (Rumelhart, McClelland & PDP Research Group 1986). But it is incorrect to equate information processing with rule-driven symbol manipulation. The point of connectionism is that a distributed system of interacting elements is able to produce behavior that was previously thought to require rules and symbols. Surely the appropriate response is to ask what makes Earle and others think that rule-driven symbolic processing is more cognitive than connectionist models. The connectionist challenge is not to cognitive models per se, but to a specific class of cognitive models based on the digital computer.

We accept Brown's concession that we did not wittingly seek to mystify or intimidate our readers. We suggest that our argument is not so confused as his quotations out of context might imply, and although we do not think we are "naïve materialists," we do take brain functioning and the constraints it places on our model seriously. Unlike Brown, we don't view our model as "merely physiological." We suspect that there is a lot packed into Brown's use of the word "merely" here — specifically, a commitment to the functionalist view that what neuroscience tells us about is irrelevant to the concerns of cognitive psychology. The underlying assumption is that physiology is irrelevant to computational issues because it is concerned only with the specifics of structural implementation (neurons, membrane constants, neurotransmitters, and so on). But this is not true. We feel that our research demonstrates that physiologists need not be (and, in fact, are not) saddled with a function-structure distinction that once and for all limits their research project to structural minutiae: Admittedly we investigated the structural properties of neural nets, but we did so in the context of viewing these networks as functional units whose input-output dynamics are captured by our model. And contrary to the functionalist

assumption, our research at the physiological level led us to produce a theory of its functional organization. We suggest that the fashionable dismissal of neuroscience popularized by a functionalism tied to the symbolic form of information processing is too simplistic in its understanding of neuroscience as practiced by researchers. (Parenthetically, our experience has been that, contrary to Brown's assertion, substantial numbers of theorists believe that the functional architecture of neural activity is identical to that of digital computers; but perhaps our dismal experience is attributable to cultural lag, and fewer people still think this way than we have inferred. Perhaps.)

What about the role of modeling, an issue raised by Brown? As a theory of whole brain function that purports to explain the changes in mass action of neurons that accompany learning new patterns of behavior, our theory spans three levels. As such it must provide the conceptual framework for the display and verification of neurophysiological correlates of behavior. And as a model of "the integrative action of the nervous system" (Sherrington 1906), it must describe or simulate the dynamic functional properties of the nervous system. Finally, because the theory attempts to explain behavior as well as brain function, it has failed if it does not yield neurobehavioral correlates or lead to methods for simulating animal and human behavior.

In previous discussions (Freeman 1981) we have adopted the tenet advanced by Craik (1952) that the essence of explanation lies in simulation. To understand some event or process, according to this view, is to generate or operate a model of it. To answer Brown, we think this can be done for our theory by using differential equations that replicate the electrophysiological patterns to make models that perform the computational operations we attribute to the brain regions we have studied. We suppose that it makes little difference a priori whether the models are made with hardware or software, although as experimentalists (naïve materialists?) we prefer to work with the former. Brown is surely as aware as we are of the pitfalls in complex programming and numerical integration; in this respect there is no advantage over hardware. In this use the computer serves as an analog, and therefore as a simile, not a metaphor. But we view the software as a helpful approach in the design of a device, just as a blueprint is a stage in the design of a tool.

We accept Barnden's criticism of our use of top-down and bottom-up; the point we wished to make is that our model resembles connectionist models in being a system that exhibits regularities and processes information without being rule-driven and manipulating symbols. Barnden also raises a physiological issue: He asks why we have not performed experiments involving the simultaneous presentation of several odors. We have not done this because it is inappropriate in our system. In olfactory physiology we have the problem of chemical reactions upon mixing odors, of new odors arising (e.g., butter and vanilla give "cake"), of optimizing the ratio of concentrations, of solubility coefficients, and so on. These factors prohibit the kind of experimentation Barnden proposes for olfaction in animals. Our theory of functional architecture should first be tested in vision or audition; if it is found to be valid, his proposed experiments can be done with relative ease.

The last word in our dialogue with psychologists we give to **Werner**. We feel that he has done a major service by expressing in clear and concise language the major theoretical implications of our target article for cognitive science. Although we originally undertook our work in order to find experimental support for the symbol-manipulationist view of information processing (Freeman 1983a), the relationship we established between measurements of behavior and electrochemical events in the nervous system forced us to adopt an alternative model (Freeman & Skarda 1985). This led us to the view that brains have a capacity to learn using cooperative activity in neural networks without anything like what the computational model based on digital computers had thought necessary. In our neural model, as in connectionist models, there is no discrete semantic interpretation given to activity in a neural net or to elements of the net; this activity varies not only with the presence and absence of particular environmental events, but also with the context. We did not come to this view without spending several frustrating years of inspecting EEG records and unit activity.

As **Werner** points out, the concepts of "representation" and "symbol" are deeply rooted in the minds of cognitive scientists; they will be eliminated or replaced only by the acquisition of a substantial body of data showing that they are unnecessary to explain behavior. Unlike **Werner**, however, we don't think that referring to these neural patterns as the "neuronal system's *own symbols*" is helpful. The distinction we wish to emphasize is not that between the first-person and third-person points of view. Our point is that the system can produce adaptive, ordered, cognitive behavior without using functional architecture based on rules and symbol manipulation. The neural patterns of our model are not symbols for the system because a distributed network doesn't require symbols to produce behavior.

Physiology: What's in a brain? The commentaries from physiologists raised several issues: the relationship of our theory to other theories, questions about the physiological role and scope of chaotic activity in the brain, and specific questions about the physiology of the system. We respond to each of these in turn.

First a word or two about the relationship of our theory to other theories: What is really new about our work? We thank **Boynton** for mentioning the pioneering work of **Demott** (1970). We have cited it in previous reports and could add to the list of pioneers in this field **Walter** (1953), **Lilly and Cherry** (1955), and **Livanov** (1977). Although these projects were able – using ingenious and imaginative devices to analyze huge quantities of data – to show that spatiotemporal patterns exist in brain activity, technical limitations prevented them from reliably reproducing or understanding the significance of the observed patterns.

Our claim to priority is not for the detection and display of patterns, nor for the technologies required for transduction, processing, and display. The electrode array is a journeyman device when seen in company with the exotic apparatus for optical, magnetic, and biochemical transduction, and our use of the digital computer is minor league in comparison with uses by meteorologists and geologists. Our primary accomplishment is the systemat-

ic measurement of repeated blocks of data, decomposition of the data into sections, the systematic testing of each of those sections for information relating to behavior, and the statistical validation of the results. By following our prescription we believe that others can find the same or similar patterns. Our priority, then, is more like that of **Columbus** than the **Vikings**: Although not the first to discover the New World, he was the first to show others how to get there and back reliably. It is this reliability of the technology to the multivariate system, and in particular the demonstration of behaviorally significant regularities in the spatial dimensions of the data, that forms the substance of our claim to priority.

Access to the digital computer is not the only basis for the difference between "seeing" that spatial patterns exist in the brain and comprehending their significance. What is crucial is the development of the necessary software. Early on we presented cinematic displays of the space-time patterns of olfactory EEG waves from electrode arrays (Freeman 1972), but the ability to demonstrate odorant specificity required that we devise the tools for measurement, a task that took 12 years to complete. Unfortunately the available theory of neural action in perception was not merely unhelpful, it was misleading. We made repeated attempts without success to locate "hot spots" of the kind purportedly revealed by 2-deoxyglucose, or to find information in phase or frequency patterns. Eventually, the key to the problem was found not in the display techniques or in the theory of nonlinear dynamics; it lay in the development of adequate techniques of measurement, including use of a behavioral assay to validate those techniques (Freeman 1987b).

We wish to point out in this connection that three technical aspects of our procedures are crucial. One is the use of spectral analyses in the temporal and spatial domains of EEG recordings as the basis for determining digitizing intervals and interelectrode distances for array recording. This, combined with theoretical analyses of the biophysical properties of the neurons generating the EEGs, provided the basis for the separation of "signal" from "noise" by filtering. A second is the use of a behavioral assay and the repetition of analytic procedures while optimizing the filter parameter for the extraction of the desired information. The third is the detailed analysis of the variance that led to our realization that the significant spatial patterns of EEG potentials were best seen after normalization of the amplitudes by channel.

Several commentators, including **Corner & Noest**, raise the issue of the relationship of our model to connectionist models. At the risk of repeating our target article, we concur with those commentators who pointed out that our approach is inherently "connectionist" and indeed "must be considered as constituting simply a possible improvement *within* that category" (Corner & Noest); our criticisms of other species within the genus should not be construed as a denial of our membership. As **Barnden** points out, our model is an "extension of present-day connectionism." We recognize that in this rapidly evolving field, some of the criticisms we directed against other connectionist models are already out of date. Nevertheless, we find that too many connectionists are preoccupied with the structural properties of their models (e.g., relationships between the numbers of nodes and

the memory capacity of a net) to the exclusion of the description and analysis of the dynamics. The tested models that we are at present acquainted with and that can be described in terms of basins and attractors are endowed with equilibrium attractors. Our results require that these be replaced with limit-cycle and chaotic attractors if they are to be relevant to the brain and the behavior it controls. In one sense this is a small step, but in many ways it is very difficult. Corner & Noest note that an extensive literature already exists on the properties of spatially distributed coupled oscillators; they sketch briefly an exciting possible route for further description and understanding of this baffling but vital system. We believe that their particular example (concerning the expression of phase patterns into spatially nonuniform amplitude patterns owing to spatial smoothing) may be directly relevant to bulbar EEG analysis, but not in the manner they suggest. As we have reported (Freeman & Baird, in press; Freeman & Viana Di Prisco 1986b), the phase pattern appears as a conic gradient in spherical coordinates, and the local differences in EEG amplitude are closely related to local differences in subsurface neural firing rates that are not subject to spatial smoothing by the volume conductor. We hope that Corner & Noest will "start delving into the complexities of structurally inhomogeneous models," and we ask them what information they need to have "specified more precisely" in order to "allow such improved models" to be developed.

In a more general vein, **Perkel** asks whether our use of nonlinear dynamics and bifurcation theory is metaphorical [the brain is (like) a . . .] or operational [the dynamics of the bulb is described by the equation $f(x) = . . .$]. We believe that it is operational, because our equations are constructed in accordance with the anatomy and known biophysical properties of component neurons and their interconnections, and are solved with boundary and initial conditions that conform to the gross anatomy and the neural input. We adjusted the parameters until the solutions to the equations conformed to the observed and measured patterns of neural activity, and we did not accept solutions for which the required parameter values were anatomically or physiologically unrealistic. In this respect our "theory" is no more or less metaphorical than any other use of descriptive equations properly selected. Superficially, at least, experimentalists can afford to be both skeptical and cavalier about theories: If a tool seems promising, we learn to use it; if it works, we continue to use it; if not, we find another. In the olfactory system we find that the language of basins and attractors helps us to assemble and simulate many aspects of patterned neural activity. The methods provide insights for further research, evidence that other parts of the cerebrum may operate in closely related ways (Freeman & van Dijk, submitted), and ways to test this hypothesis. What more can one ask of theory?

Babloyantz points out that ours isn't the first physiological account to postulate chaotic activity based on EEG recordings. She provides a useful list of references to recent work on the dimensional analysis of putative chaos in human scalp EEG recordings, to which we add work by Nicolis (1985b) and Nicolis and Tsuda (1985) on chaotic dynamics of information processing by the brain. (**Perkel**, by the way, asked whether chaos is unique to olfaction. Babloyantz and others have clearly demonstrated that it is

not.) These and related studies have established that low dimensions appear in analyses of records from subjects in deep sleep and in certain forms of epilepsy. But in these studies the estimated dimensions of waking EEGs are so high that the distinction between chaos and noise or a mix of the two becomes blurred (see commentary by **Thom**). A further difficulty with these studies is that the single channel of the scalp EEG is undefined with respect to the numbers of functional entities (and therefore of dimensions) that contribute to the record (as distinct from our deliberate restriction to activity from a single entity); and the scalp EEG is subject to much stronger spatial and temporal smoothing than pial recordings, leading to artifactual reduction in the apparent dimension. So, although our research is not the first to postulate chaotic activity based on EEG recordings, we do believe our work makes important new contributions by eliminating some of these difficulties.

Contrary to **Babloyantz's** assertion, our main point is not that brain activity conforms to the dynamics of chaos, but that the brain organizes its own space-time patterns of function and thereby its own structure. We postulate that it generates chaotic activity as an essential precursor to the emergence of ordered states. Far from being misleading, we think that our statement that "chaos is a controlled noise" is appropriate. The patterns of activity we observe in the bulb have commonality of waveform over cortical regions comprising hundreds of millions of neurons, with phase gradients and spectral distributions (both temporal and spatial) that are held within narrow limits; amplitudes (root mean square) that are regulated precisely in accordance with motivational state; and, above all, the maintenance of sustained basal activity without recurring spatial patterns. The generator of this activity is exceedingly robust, a neural mechanism that we can readily believe has been present in vertebrates for over four hundred million years. The activity looks like "noise," serves (we believe) purposes met by unstructured or pseudorandom activity, and is turned on and off rapidly and reliably with respiration in a controlled manner. We agree with Babloyantz that to understand our target article properly readers will have to seek out and study "more technical publications" as cited here and elsewhere, and we hope that the paper will motivate some of them to do so.

As to our view of the underlying physiology, **Perkel** questions why we propose that strengthening of excitatory synapses, but not of inhibitory ones, occurs with learning. We base our hypothesis on measurements of changes in the waveform of averaged evoked potentials when animals are trained to respond to electrical stimuli (Emery & Freeman 1969), and on determinations of the parameter changes that are necessary and sufficient to replicate these pattern changes in the solutions of differential equations that model the neural dynamics (Freeman 1979b). Interestingly, several theoretical advantages accrue from the experimental result that only the excitatory synapses change. One is the exquisite sensitivity of the olfactory system, arising from the form of the sigmoid curve under recurrent excitation. Another is the pattern stabilization and figure completion that results from strengthened excitatory connections (Babloyantz & Kaczmarek 1981). Yet another is the exploitation of the Hebb (1949) rule, which is the basis for learning under rein-

forcement in our model; it is difficult to see how this rule might be implemented toward formation of a nerve-cell assembly if the synapses to be strengthened were inhibitory. The findings of Wilson, Sullivan, and Leon (1985) are fully consistent with our model; we have shown elsewhere (Gonzalez-Estrada & Freeman 1980) that observed suppression of mitral cell discharge can be the manifestation of profound excitatory action onto those same cells, given the proper system parameters. Here, indeed, is an opportunity for theory to come forward and explain the paradoxical and counterintuitive.

Perkel also raises what he refers to as the "serious problem" of how the system can read out the identity of a familiar odorant when static and invariant representations do not exist in the bulb or elsewhere in the olfactory system as the basis for olfactory memory. We have proposed elsewhere (Freeman & Skarda 1985) that the coherent, phase-locked activity generated by mitral cells, falling onto the prepyriform cortex after spatial and temporal reorganization in the olfactory tract, causes further bifurcation in that structure, initiating the process of response selection. But Perkel's question is important, because we do not yet know how that is done. For us, the really serious issue is not that the event is a transitory bundle of energy rather than a fixed state; it is that of developing a model based on distributed dynamic networks that can explain how a dynamic state (which need not have gone to completion) can lead to state changes in the rest of the nervous system leading to a response. One advantage of a connectionist model is that it doesn't require the fixed states that symbolic processing requires, and it allows us to conceive of new forms of interaction among subsystems like those found in the brain.

Corner & Noest's chief challenge to our target article concerns the issue of projection of afferent activity from the receptors to the bulb. As they see it, each odorant stimulus activates a set of receptors that in turn activates its odorant-specific pattern in the bulb, with "magnification of preexisting differences in the spatial distribution of afferent signals. . . . What, then, needs to be 'learned' about such signals? . . . Nothing else, surely, than . . . behavioral significance." Corner & Noest fail to grasp that the major task for learning to identify an odor is the formation of a nerve-cell assembly by the pair-wise strengthening of synapses between co-activated mitral cells (Hebb 1949). This task reflects the necessity for establishing an equivalence over all receptors (and the mitral cells to which they transmit) that are sensitive to a particular odor in an invariant manner. This is the key problem that Lashley (1942) posed in terms of stimulus equivalence. The formation of the nerve-cell assembly, in accordance with the postulate proposed by Hebb, takes place only under reinforcement and involves the release of norepinephrine into the bulb (Gray, Freeman & Skinner 1986). As we discussed in our target article, the EEG spatial patterns are as closely related statistically to the CR (conditional response) as to the CS (conditional stimulus). In answer to Corner & Noest's specific query, there were nine distinctive patterns for each of the four subjects that learned the discriminations, including discriminable spatial patterns that had the same odorant as CS but with different "meaning" (that is, that elicited a different CR) at different stages in the training program.

Finally, Corner & Noest mention several issues that, although not relevant to the main points raised in our target article, are of special interest to physiologists, so we'll address them here. First, we think that the resemblance Corner & Noest point out between the classical brainstem mechanisms for cortical arousal, on the one hand, and the energizing effect of receptor input to the olfactory bulb, on the other, is not to be taken seriously. The nonspecific arousal process is mediated by ascending reticular axons operating on and through the thalamic reticular nuclei, whereas specific sensory activity is carried by distinctive sensory pathways through modality-specific thalamic nuclei. The "dual" input to the bulb is carried by one and the same pathway; receptor axons carry action potentials that bear specific information by depolarizing the apical dendrites of selected mitral cells. Massive depolarization of the whole system brings about a bifurcation that leads to response selection. This is not the "preparation" of the bulb by the prior action of a parallel ascending pathway; it is a result of the concomitant induction of an instability of the system receiving the input. Moreover, the time and distance scales of these phenomena are significantly different. Reticular activation is broadcast to all parts of the nervous system by ascending and descending projections, irrespective of the modality of the arousing stimulus, and the aroused state tends to last for seconds to minutes with gradual abatement. The transition in the olfactory system is localized to the bulb and cortex, and it terminates during exhalation in a fraction of a second. Finally, the arousal response is centrifugally induced in the olfactory system as well as in other sensory cortices, whereas the formation of the burst is dependent on the centripetal sensory input and requires that the bulb already be in the aroused state.

Second, we are in fact unable to explain alpha suppression in arousal or the enhancement of hippocampal theta in certain states involving orienting, but these phenomena lie outside the scope of our data and models. Our models do not generate activity in the alpha and high theta ranges (roughly 5 to 15 Hz) without impermissible parameter settings, but neither does the olfactory system.

Third, it is not the case, as stated by Corner & Noest, that gamma EEG waves "have gone undetected" in the neocortex. Systematic studies as well as anecdotal reports abound attesting to the presence of "40 Hz" activity, as it is commonly called, in many areas of the neocortex (e.g., Chatrian, Bickford & Uihlein 1960; Sheer 1976). In our opinion Corner & Noest underestimate the strength of the signal degradation imposed on neocortical EEG potentials by the spatial dispersion of the generating cells in directions perpendicular to the pia. The three cortical structures with the most striking alignment of their generating cells in this respect are the bulb, the prepyriform cortex, and the hippocampus, and these three have EEG amplitudes that easily exceed the amplitudes of neocortical EEGs by 10- to 20-fold. Moreover, the relatively low amplitudes of that activity, especially at the scalp, are confounded by electromyographic (EMG) potentials that badly obscure gamma activity (40 to 90 Hz). Hence the gamma activity is known to exist, but it is poorly documented and has largely been ignored.

Fourth, we agree with Corner & Noest that "the basic

notion of widely synchronized neuronal carrier waves . . . is by no means excluded by differences among brain regions displaying the precise [frequency] characteristics of these (chaotic) waves," but we do not consider as "major" the challenge to account for the observed frequencies and amplitudes of olfactory EEG activity or to show their theoretical advantages. We have demonstrated repeatedly that the gamma range is the characteristic frequency band for neurons with passive membrane time constants on the order of 5 msec, and that the amplitudes are the result of the cytoarchitecture of the laminar structures. These are as they are because of the properties of the neurons that comprise the areas and generate the waves, not because these properties critically influence olfactory information processing. We suspect that neocortical cell assemblies tend much more strongly to chaotic activity, which renders them all the less accessible to our present understanding; but we do not consider it incumbent on us at this time to explain why this is so, or to show what advantages might accrue to vision or audition thereby. We have attempted to understand and explain paleocortical dynamics and to speculate a bit about the neocortex, not to propose a general theory of the EEG.

Mathematics: The uses and abuses of chaos. Chaos and its possible role in pattern recognition figured prominently in our target article. We suggested that in the brain chaos is necessary for learning new odors. Not all the commentators agreed with this, but as Garfinkel noted, the view that chaos plays a functional role is "something of an about face" for a phenomenon that has traditionally been viewed as highly undesirable. We are especially encouraged by the examples he cites that indicate chaos is not merely tolerated but essential for optimal performance of systems in search of their own goals or states of minimal energy. These uses for chaos translate readily into the maintenance of background activity while avoiding hypersynchrony, as in epilepsy; the flexibility and adaptiveness of behavior in the face of unpredictable environments; and the speed of operation of brains in entering and leaving states sequentially. We suggested that flexibility in responding to the changing olfactory environment is provided by the chaotic basal state, and that chaos doesn't merely provide noise in the manner of a Boltzmann machine to avoid local minima in a convergence process, but that it allows a relatively high energy state to be maintained between signal episodes, so that the neural system does not have to be dragged out of or dropped into a deep energy well with each bifurcation. It can flip lightly and quickly with each sample and flip back again. Furthermore, because the same mechanism generates both chaos and carrier, the "noise" is shut off when the "signal" goes on and vice versa. The signal is detected between the noise periods and not in them, so that the "signal/noise ratio" concept is not applicable here.

Most of the objections to our use of chaos were based on proposed alternative models that don't require chaotic activity. The commentaries of Grossberg, Levine, and Rosenfeld, Touretzky & the Boltzmann Group belong to this category. By way of preface, however, we want to make our general position on chaos clear. We are the first

to admit that we have no proof that chaos is essential to pattern recognition, whether biological or artificial, or that nonzero equilibria under noise might not serve as well. We believe, however, that we have shown that chaos exists in the olfactory system, and that our suggestions as to its roles are plausible and useful; certainly chaos should not be averaged out, discarded, or ignored. Although our understanding of chaos is rudimentary in comparison with our needs, the most effective way to proceed is by close cooperation between theorists and experimentalists, as exemplified by our exchanges with Garfinkel, especially in the analysis of spatially distributed systems of coupled oscillators. Now, on to some of the objections.

In the past two decades Grossberg and his associates have consistently produced imaginative, detailed, yet comprehensive models expressing the formal bases of learned behavior generated by nervous systems. We note, however, that Grossberg's "*Gedanken*" experiments have been designed primarily to explain phenomena deriving from psychophysics; the relationship to neurophysiology occurs through his use of neural "metaphors," such as the cellular dipole representing local inhibitory feedback, the first-order decay process representing passive membrane, shunting inhibition, and the modifiable synapse. We view these terms as metaphorical because, in using them, Grossberg normalizes his state variables to dimensionlessness in time and space. In principle, of course, he could retain the conversion factors and return to the metric of the relevant nervous system, but in practice he does not because his avowed interest lies in general principles and not in specific examples. It is accordingly necessary for experimentalists to supply the conversion factors in order to test his theories, to the extent that they are intended to explain the brain.

Grossberg's claim, based on his work on olfactory coding and its explanation by his adaptive resonance model (1976; 1981) and recently incorporated in his ART 1, is that chaos is unnecessary in his model. What are we to make of this claim in relation to our data? A comparison of ART 1 with our KII model shows that both consist of excitatory and inhibitory neurons formed into three or more serial layers with massively parallel axonal connections between them. The initial layer in both models is comprised of the sensory transducer neurons; Grossberg's layer S1 purports to embody our olfactory bulb, and his layer S2 the prepyriform cortex. His element for arousal from mismatch or attentional biasing may correspond topologically to the anterior olfactory nucleus, because this is a key site for centrifugal brainstem control of the bulb. Within each layer there is local negative (inhibitory) feedback, and there is extensive feedback between layers S1 and S2. Both systems have the ability to modify synaptic weights under reinforcement during learning. Both invoke the sigmoid curve as the static nonlinearity that dominates the dynamics.

However, despite these superficial resemblances the differences between the two systems are so great that direct comparisons tend to be nonproductive and misleading. In ART 1 the most important modifiable synapses are at the input from layer S1 to layer S2, equivalent to that between the lateral olfactory tract and the pre-

pyriform pyramidal cell, and from S2 to S1 equivalent to that between the medial olfactory tract and the granule cell (not the mitral cell, as Grossberg's model states). In ART 1 the input pattern is sustained and maintained for matching purposes at the input synapse to S1, equivalent to the primary olfactory nerve synapsing on mitral cells. In our KIII set the input is ignored after bifurcation. Moreover, this synapse is the site of dynamic range compression and signal normalization (presynaptic inhibition); none of these synapses change in relation to associative learning in KIII, though they are subject to posttetanic and long-term potentiation. Shunting inhibition in ART 1 is multiplicative; recurrent inhibition in KIII is additive. In ART 1 local inhibitory feedback serves primarily for contrast enhancement. In the KIII set it serves primarily for generating the carrier frequency of bulbar output. In Grossberg dynamics, the state variables tend to fixed values (equilibrium attractors). Long-range excitatory connections are not prominent in ART 1; the distance of excitatory transmission is kept below that of inhibitory transmission. In the KIII set long-range excitatory connections in both the bulb and cortex are crucial for the formation of Hebb-type nerve-cell assemblies, because it is here that the synapses are modified in associative learning. In ART 1 the Hebb rule is applied to modification at the input level by virtue of feedback of spatially detailed and precisely timed information from S2 to S1. This cannot occur in the KIII set, because it models the feedback from the prepyriform cortex to the bulb, and this has such marked spatial divergence and temporal dispersion in both forward and feedback directions that no such transmission of detailed information is feasible. The pathways act as strong low-pass filters and remove it.

From these and related anatomical and physiological measurements we conclude that function in the olfactory system cannot depend on precise timing and precise topographic mapping. Its algorithms must be reliable in the face of continual smoothing of activity by temporal dispersion under axonal transmission and by spatiotemporal integration in dendrites. This is one of several reasons that we insist repeatedly that behaviorally relevant neural information is to be found in the average activity of ensembles (as manifested in the EEG) and not in the activity of single neurons, once the first stage of carriage by sensory neurons has been passed.

In addition, although Grossberg has often stated that his models hold good for both the single neuron and for the ensemble, we maintain that this cannot be so, because the static sigmoid nonlinearity that dominates the dynamics of both ART 1 and the KIII model holds only for the ensemble and not for the single neuron. He has also claimed that his results hold for both steady-state and oscillatory solutions (equilibrium and limit-cycle attractors). In our view he has not adequately demonstrated the pattern-recognition dynamics of the oscillatory standing-wave type of system to substantiate this claim.

From these considerations it should be apparent that although we and Grossberg appear to be thinking and writing about the same nervous system, in actuality we are skew, almost entirely disjunct, because of the differences between our methods, values, and data bases. Of course, Grossberg's assertion would be valid if we had claimed that chaos was required for all pattern-recognition

devices. Such a claim, however, would be foolish. We readily acknowledge the validity of his claims about equilibrium solutions for ART 1, but we want to point out that his dynamics have not yet been developed sufficiently to simulate the actual performance of our KIII model. Until he is able to expand his "minimal anatomies" to include and fully exploit modifiable excitatory crosscoupling within S1 and within S2, and to establish the proofs of performance of this system with periodic attractors, as he has for the fixed points of ART 1, we cannot concede that he has a counterexample. His model is not in the same domain. We conclude that Grossberg's cautionary note that chaos is inessential in ART 1 is important to consider, but that we must await further development of his ART models that explicitly exhibit the anatomy and dynamics we observe in the brain before we can be bound by his logic. And should his models prove to be less flexible than he might desire, we recommend a small dose of chaos.

In a related vein, we accept Levine's point that some connectionist models have the capacity to respond to novel as well as to familiar input, so that input can be rapidly identified as novel and the learning process can begin. Such a process is embodied in Grossberg's ART 1 model and in work cited by Levine. Our difficulty, as discussed in our response to Grossberg above, is that our analysis indicates the olfactory system is not capable of performing the operations required for match and mismatch detection, at least in the manner carried out by ART 1. This is because the feedback pathway from the putative layer S2 (the prepyriform cortex) to layer S1 (the olfactory bulb) is incapable of sustaining transmission of information with the requisite specificity of timing and spatial resolution. That is, we see no realistic way that fibers in the medial olfactory tract can return an organized pattern from the prepyriform cortex to the bulb and "match," "correlate," or "compare" it with a pattern in the bulb that is sustained by input from the receptors. Furthermore, our recordings tell us that after bifurcation the bulbar activity pattern reflects the generalization to a stereotypic form affected by a class of inputs based on experience rather than current individual inputs. We therefore contend that although Levine and Grossberg are on the track of devices that may far outperform the simpler connectionist models currently in vogue, these do not help to explain the dynamics of those parts of the nervous system with which we are familiar. If their models can be made to do what brains can do, or even to outperform them, then they need not trouble with chaos. If they cannot, then, like theorists of the past several decades, they might consider returning to the nervous system for some more insights and ideas. This is what von Neumann (1958) did in his quest for mastery of the newly conceived programmable digital computer.

The other large-scale assault on our target article from connectionists came from Rosenfeld et al. who presented clear-cut arguments strongly defending their views and practices without effectively, we think, eroding our claims or responding to our findings or our speculations concerning the significance of their models for cognitive science. We recognize and have repeatedly stated that all models are abstractions, that the value of each model depends on how well the selection of detail to be included or excluded helps the modeler to attain a stated goal, and

that judicious selection is more to be valued than mere fidelity to some "view" of how an aspect of the brain being modeled actually works. That "view" in itself being a hypothesis and therefore a model of sorts, we think that Rosenfeld et al. have invited us into a cul-de-sac with their allegation that we are confused. Furthermore, we believe that they have short-changed themselves in their efforts to "contribute to the emergence of a new taxonomy" of connectionist models by their apparent refusal to add some more cages to their zoo for models with dynamic attractors. To be sure, we have no evidence for the computational superiority of dynamic attractors in mathematical form, nor do we have bench-mark studies comparing the speed, accuracy, capacity, stability, and so on of two or more models having point versus chaotic attractors that perform some common task. We are not aware that benchmark studies exist as yet for comparative evaluation of different species of point models. But it is not true that we have no "evidence that could support [our] claim" for the value of dynamic attractors. The bulk of our target article presents evidence that the olfactory system works with dynamic attractors; we cited additional evidence that suggests the rest of the cerebrum may do so as well. This type of evidence is considered by some connectionists to be relevant to artificial intelligence (AI), considering that the performance of the brain still outclasses those of AI models.

Rosenfeld et al. ask us, "Which of the characteristics of chaos are necessary to the role it plays in generating new attractors, and which are irrelevant?" Our answer is that the following characteristics are necessary. We have stated that in order for the Hebb rule to operate, the neurons involved in the learning process must be active during the CS-induced activity under reinforcement of the UCS (unconditional stimulus), but that the activity must not conform to the attractor of a known odor. Hence, if the intensity of the activity must be high and its pattern should not conform to any previous spatial pattern of activity, then the pattern must appear to be random – i.e., chaotic – and must cover the entire bulb. On the other hand, the following aspects are irrelevant. The time series looks random. We find that its spectrum is broad in comparison to the spectrum of bursts with known odorants, and there is more low temporal frequency energy.

Rosenfeld et al. write of a "forced reset action" existing in the dynamics of the bulb during exhalation, that is analogous to the action of turning their devices "off." This is a significant mischaracterization of the dynamics of our model. During inhalation there is a forced choice leading to capture of the system at a "single (albeit dynamic) state that is computationally equivalent to the corner of a hypercube" (hypertorus?), which in our model is due to an obligatory side effect of the input surge; but during exhalation there is *relaxation* to the basal state, not forced reset. Furthermore, there is no reset to the "center of the hypercube," even if the center could be computed within a reasonable time; instead, in our model *the entire hypercube vanishes*. It is re-created with each new inhalation causing bifurcation; this is the essence of self-organization.

Rosenfeld et al. have likewise failed to grasp the significance of our remarks on "pattern completion." There is a simple sense of "pattern completion" that can

be said to take place: the conjectured extensive spread of activity through a nerve cell assembly (NCA) when any subset of its neurons is excited. There is another sense that is untenable: This is to suppose that the NCA is like the form of the letter "A" that is filled in by mutual excitation within the NCA following excitation of some fraction of its parts. The NCA is formed by repeated presentations of an odorant at sufficiently dilute concentration to prevent adaptation, and we suppose that the entirety is never activated on any one presentation, nor would that presentation be in any sense necessary, crucial, or identifiable even if it did occur. Furthermore, the successful convergence into the basin of a correct attractor, according to our model, depends on activating an NCA but does not specify that the activation need be complete, or even need asymptotically approach completion. The outcome may be a behavioral pattern, so that a stimulus-response configuration may be said to go to completion, but this does not require that an internal dynamic activity pattern go to "completion" in some kind of exemplary archetypal or normative state. The regularities in spatial pattern we have observed do represent possible outcomes, but cannot be shown, we believe, to be "completed" in the sense commonly used. We agree that our notion of "destabilization" is a metaphor until realized in software or hardware, and we intend to pursue it further; but we deny that generalization over equivalent stimuli, which is done by a basin and its attractor, is the same as pattern completion.

Moreover, in response to Rosenfeld et al. we wish to point out that we cannot deal with the concept of "similarity" in our model for the same reason that we cannot handle "concentration." Both of these require serial processing of successive images or samples with comparison over time. We emphasize again that our model deals only with preattentive [or what Julesz (1984) terms "pop-out"] cognition. We do not have anything useful to say about attentive cognition, except that in our opinion it must involve proprioceptive and reafferent information so that successive sensory information samples can be combined with the information about what is done to get them. We also think it unlikely that each modality will be found to have such neural machinery separately, so that it should be sought after the combination of sensory input from all modalities into gestalts. From neurological considerations the most likely site of convergence is the entorhinal cortex (Lorente de Nó 1934), for which the hippocampus may serve as a stack register for temporal integration of serial gestalts. Much work needs to be done on multiple coexisting NCAs, but not in the context chosen by Rosenfeld et al. or at least not in neurobiological studies. As a postscript, we share their aversion to blind simulation of neural circuitry in the absence of any analytic handle on it, which is our reason for having made such a heavy investment in linear analysis of neural dynamics (Freeman 1972; 1975; 1987b). We would also like to agree on "minimal assumptions" and "essential features," if only we knew them in advance.

Finally, while we have been occupied with adding some neural-based variations to the kinds of models contained in the connectionist zoo, and the connectionists were busy questioning our reasons for doing so, Thom was taking a shot at us all. We understand Thom's suspicion that at present "connectionist models' ulti-

mately mean very little," especially those deriving from the spinglass analogy. However, we do not understand what he means when he says that "a neural mass exhibits internal symmetry of a geometric type (translation, rotation, etc.)." We assume that the collection of local elements comprising a mutually excitatory set is interconnected by feedback connections between each pair of elements, and that the gain coefficients representing synaptic weights can take the form of a matrix. In the naïve state we assume that the weights are uniform or nearly so. With learning under the Hebb rule, in which the connections between pairs of neurons are more strongly weighted by co-activity, the symmetry of the matrix in our model is preserved, because each pair of neurons is reciprocally connected. The weights in both directions depend on the correlation of activity by the same pair. Symmetry does not hold for the negative feedback connections between excitatory and inhibitory cells, but in our model these weights are fixed and can be partialled out. Recent advances have shown that symmetry in this sense need not be maintained in connectionist models, and that with sufficient asymmetry limit-cycle attractors appear. Surely this is not surprising, but these more complicated systems are more difficult to comprehend.

We are also puzzled by Thom's statement that "such chaotic systems have no equations." If by this he means that the EEG cannot be simulated or fitted by a closed function such as cosine or Bessel, we surely agree. But we have shown that high-dimensional sets of coupled ordinary differential equations (ODEs) generate activity that is statistically indistinguishable from the EEG. Why is it permissible to call it chaos before we find the ODEs, but not after? We suspect that Thom's criticisms of terminology are directed more toward his fellow mathematicians than toward us. Certainly there are degrees of unpredictability and sensitivity to initial conditions, ranging from a barely detectable wobble about a point or limit-cycle trajectory to the sort of wildness that "cannot be explicitly described." Our ODEs simulate all of these gradations.

We are willing to adopt whatever convention mathematicians eventually agree upon. Thom offers the choice between an oxymoron and a neologism (between "self-organizing process" and "chreod") to label the process of the emergence of order in a system without prior specification from the outside. What is important to us is not the name, which explains little, but the concept by which we conceive of a very large number of neurons that are coupled into a coherent mass, which, when highly interactive, has degrees of freedom far lower than the number of neurons, perhaps so few as can be counted on the fingers of one hand or two. The finding that ODEs modeled after the olfactory system can be made to generate olfactory seizure patterns as unpredictable as one finds in nature is for us a liberation from the tyranny of Fourier decomposition. We recognize the weakness of the measures of the invariants of chaos that mathematicians have thus far made available to us (see our reply to **Babloyantz**), but that problem is not germane to questions of terminology. It is obvious that the whole field of dynamical systems is uncomfortable with terminology in transition. Perhaps it is a measure of our immersion in reflex determinism that we have such difficulty finding

words and concepts for these common phenomena. Thom compounds our discomfort by his final comment that "Pavlovian conditioning . . . could be explained as a purely dynamical effect" (if only he had stopped here) "of resonance"! The bulk of our target article, and for that matter of most of connectionism, is devoted to explaining the energy-consuming dynamical character of brain function, but the term "resonance" – as in sympathetic vibrations or a Helmholtz resonator, the passive transfer or accumulation of energy at specific temporal frequencies, the ringing of tuned oscillators – is empty at best and misleading and obfuscating at worst. We urge that the language of basins and attractors be used; perhaps this will also have to be discarded at some future time when it has also been debased by the verbal inflation that comes with overuse, but surely by then we will have a new vocabulary to debate with and about.

A final point needs to be made before we conclude. Thom is too generous in characterizing our experimental data as having "gaps"; at best, they constitute a small clearing in a large forest. The particular experiment he describes, in which one odorant serves as a CS for an appetitive UCS at another stage, has shown that the two EEG patterns differ. More to the point, when a rabbit is conditioned aversively to respond seriatim to odorant A, then B, C, and D, and is again conditioned to A, the spatial pattern changes with each new odorant, but it does not revert with reintroduction of odorant A to pattern A on the first conditioning. It changes to a new pattern on the repeat conditioning (Freeman & Schneider 1982). This result is contrary to expectations based on analogy with digital computer memory. It says something profound not so much about brains as about our preconceptions about the static nature of memory and our need to believe in mnemonic invariants. One might ask, how could the rabbit know that it was the same odorant in both stages of conditioning if there was a different spatial pattern in the second stage? The answer is that we have no sure way of knowing whether the animal can retain such information and make such judgments over the weeks required to do the experiment, and even if it could, which seems highly unlikely, the background and context are changed, and these are influential in shaping the forms of associative memories.

We do not commit Whitehead's (1960) "fallacy of misplaced concreteness." We do not claim an infinite storage capacity in the bulb for osmic memories; on the contrary, the psychophysical studies referred to in our target article show that the retentive capacity of the olfactory analyzer, if language is not used, is limited to about 16 odors at any given time. But the learning ability allows the repertoire to be modified and updated without increasing the total content. We have not yet attempted to address the questions of how new NCAs are overlaid or intertwined with preexisting ones, or what happens to the old ones, how they are deselected, and so forth. This is fertile ground for further studies in physiology, theory, and hardware modeling.

References

- Abraham, R. H. & Shaw, C. D. (1985) *Dynamics, the geometry of behavior*, vol. 2. Ariel Press. [aCAS]

- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science* 9:147-69. [RR]
- Anderson, J. A., Silverstein, J. W., Ritz, S. A. & Jones, R. S. (1977) Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review* 84:413-51. [aCAS]
- Arbib, M. A. (1972) *The metaphorical brain*. Wiley-Interscience. [DHP]
- Ashby, W. R. (1952) *Design for a brain*. Chapman & Hall. [aCAS, DHP]
- Auslander, D., Guckenheimer, J. & Oster, G. (1978) Random evolutionarily stable strategies. *Theoretical Population Biology* 13:276-93. [AG]
- Babloyantz, A. & Destexhe, A. (1986) Low-dimensional chaos in an instance of epilepsy. *Proceedings of the National Academy of Sciences of the United States of America* 83:3513-17. [a CAS, AB]
- (in press) Strange attractors in the human cortex. In: *Temporal disorder and human oscillatory systems*, ed. L. Rensing, U. an der Heiden & M. C. Mackey. Springer. [AB]
- Babloyantz, A. & Kaczmarek, L. K. (1981) Self-organization in biological systems with multiple cellular contacts. *Bulletin of Mathematical Biology* 41:193-201. [rCAS]
- Babloyantz, A., Nicolis, C. & Salazar, M. (1985) Evidence of chaotic dynamics of brain activity during the sleep cycle. *Physics Letters* 111A:152. [AB]
- Baird, B. (in press) Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. *Physica D*. [aCAS]
- Barlow, H. B. (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1:371-94. [aCAS]
- Bénard, H. (1900) Les tourbillons cellulaires dans une nappe liquide. *Revue Générale des Sciences Pures et Appliquées* 12:1309-28. [RT]
- Bressler, S. (1987a) Relation of olfactory bulb and cortex I: Spatial variation of bulbo-cortical interdependence. *Brain Research*. 409:285-93. [aCAS]
- (1987b) Relation of olfactory bulb and cortex II: Model for driving of cortex by bulb. *Brain Research*. 409:294-301. [aCAS]
- Brown, T. H., Perkel, D. H. & Feldman, M. W. (1976) Evoked neurotransmitter release: Statistical effects of nonuniformity and nonstationarity. *Proceedings of the National Academy of Sciences of the United States of America* 73:2913-17. [DHP]
- Bullock, T. H. & Horridge, G. A. (1965) *Structure and function in the nervous systems of invertebrates*, vol. 1. W. H. Freeman. [aCAS]
- Burns, B. D. (1958) *The mammalian cerebral cortex*. Edward Arnold. [aCAS]
- Cain, W. S. (1980) Chemosensation and cognition. In: *Olfaction and taste VII*, ed. H. van der Starre. IRL Press. [aCAS]
- Carpenter, G. A. & Grossberg, S. (1987) A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing* 37:54-115. [SG]
- Chatrjian, G. E., Bickford, R. C. & Uihlein, A. (1960) Depth electrographic study of a fast rhythm evoked from human calcarine region by steady illumination. *Electroencephalography and Clinical Neurophysiology* 12:167-76. [rCAS]
- Chay, T. R. & Rinzel, J. (1985) Bursting, beating, and chaos in an excitable membrane model. *Biophysical Journal* 47:357-66. [DHP]
- Churchland, P. S. (1980) A perspective on mind-brain research. *Journal of Philosophy* 77:185-207. [RB]
- (1986) *Neurophilosophy: Toward a unified understanding of the mindbrain*. MIT/Bradford. [aCAS]
- Cohen, M. A. & Grossberg, S. (1983) Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Transactions on Systems, Man, and Cybernetics* 13:815-26. [DSL]
- (1986) Neural dynamics of speech and language coding: Developmental programs, perceptual grouping, and competition for short term memory. *Human Neurobiology* 5:1-22. [SG]
- Conrad, M. (1986) What is the use of chaos? In: *Chaos*, ed. A. V. Holden. Manchester University Press. [aCAS]
- Cowan, J. D. (1968) Statistical mechanics of nerve nets. In: *Neural networks*, ed. E. R. Caianiello. Springer-Verlag. [DHP]
- Craik, K. (1952) *The nature of explanation*. Cambridge University Press. [rCAS]
- del Castillo, J. & Katz, B. (1954) Quantal components of the end-plate potential. *Journal of Physiology (London)* 124:560-73. [DHP]
- DeMott, D. W. (1970) *Toposcopic studies of learning*. Charles C. Thomas. [rCAS, RMB]
- Dretske, F. (1986) *Boston studies in the philosophy of science*. Vol. 90: *Minds, machines and meaning, philosophy and technology II*, ed. C. Mitcham & H. Hunning. Reidel. [GW]
- Dreyfus, H. L. (1972) *What computers can't do*. Harper & Row. [RB]
- Dreyfus, H. & Dreyfus, S. (1986) *Mind over machine*. Free Press. [aCAS]
- Ellias, S. A. & Grossberg, S. (1975) Pattern formation, contrast control, and oscillations in the short term memory of shunting on-center off-surround networks. *Biological Cybernetics* 20:69-98. [SG]
- Emery, J. D. & Freeman, W. J. (1969) Pattern analysis of cortical evoked potential parameters during attention changes. *Physiology and Behavior* 4:60-77. [rCAS]
- Ermentrout, B., Campbell, J. & Oster, G. (1986) A model for shell patterns based on neural activity. *Veliger* 28(4):369-88. [aCAS]
- Farmer, D., Hart, J. & Weidman, P. (1982) A phase space analysis of baroclinic flow. *Physics Letters* 91A:22-24. [AG]
- Feldman, J. A. & Ballard, D. H. (1982) Connectionist models and their properties. *Cognitive Science* 6:205-54. [aCAS]
- Feyerabend, P. (1978) *Against method*. Verso. [CW]
- Fienberg, S. E. (1974) Stochastic models for single neurone firing trains: A survey. *Biometrics* 30:399-427. [DHP]
- Fodor, J. A. (1983) *The modularity of mind: An essay on faculty psychology*. Bradford/MIT Press. [rCAS]
- Freeman, W. J. (1972) Waves, pulses and the theory of neural masses. *Progress in Theoretical Biology* 2:87-165. [arCAS]
- (1975) *Mass action in the nervous system*. Academic Press. [arCAS]
- (1978) Spatial properties of an EEG event in the olfactory bulb and cortex. *Electroencephalography and Clinical Neurophysiology-EEG Journal* 44:586-605. [aCAS]
- (1979a) Nonlinear gain mediating cortical stimulus-response relations. *Biological Cybernetics* 33:237-47. [aCAS]
- (1979b) Nonlinear dynamics of pleocortex manifested in the olfactory EEG. *Biological Cybernetics* 35:21-34. [arCAS]
- (1979c) EEG analysis gives model of neuronal template-matching mechanism for sensory search with olfactory bulb. *Biological Cybernetics* 35:221-34. [aCAS]
- (1980) Use of spatial deconvolution to compensate for distortion of EEG by volume conduction. *IEEE Transactions on Biomedical Engineering* 27:421-29. [aCAS]
- (1981) A physiological hypothesis of perception. *Perspectives in Biology and Medicine* 24:561-92. [arCAS]
- (1983a) Physiological basis of mental images. *Biological Psychiatry* 18:1107-25. [arCAS]
- (1983b) Dynamics of image formation by nerve cell assemblies. In: *Synergetics of the brain*, ed. E. Basar, H. Flohr & A. Mandell. Springer-Verlag. [aCAS]
- (1986) Petit mal seizure spikes in olfactory bulb and cortex caused by runaway inhibition after exhaustion of excitation. *Brain Research Reviews* 11:259-84. [aCAS]
- (1987a) Simulation of chaotic EEG patterns with a dynamic model of the olfactory system. *Biological Cybernetics* 56:139-50. [aCAS]
- (1987b) Techniques used in the search for the physiological basis of the EEG. In: *Handbook of electroencephalography and clinical neurophysiology*, vol. 3A, part 2, ch. 18, ed. A. Cevins & A. Remond. Elsevier. [arCAS]
- Freeman, W. J. & Ahn, S. M. (1976) Spatial and temporal characteristic frequencies of interactive neural masses. *Proceedings, Institute of Electronics and Electrical Engineering, International Conference on Proc. IEEE. Intern. Conf. Cybernetics and Society* 1-3:279-84 [aCAS]
- Freeman, W. J. & Baird, B. (in press) Correlation of olfactory EEG with behavior: Spatial analysis. *Behavioral Neuroscience*. [aCAS]
- Freeman, W. J. & Grajski, K. A. (in press) Correlation of olfactory EEG with behavior: Factor analysis. *Behavioral Neuroscience*. [aCAS]
- Freeman, W. J. & Schneider, W. S. (1982) Changes in spatial patterns of rabbit olfactory EEG with conditioning to odors. *Psychophysiology* 19:44-56. [arCAS]
- Freeman, W. J. & Skarda, C. A. (1985) Spatial EEG patterns, nonlinear dynamics and perception: The neo-Sherringtonian view. *Brain Research Reviews* 10:147-75. [arCAS]
- Freeman, W. J. & van Dijk, B. (submitted) Spatial patterns of visual cortical fast EEG during conditioned reflex in a rhesus monkey. [aCAS]
- Freeman, W. J. & Viana Di Prisco, G. (1986a) 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. [aCAS]
- (1986b) Correlation of olfactory EEG with behavior: Time series analysis. *Behavioral Neuroscience* 100:753-63. [arCAS]
- Froehling, H., Crutchfield, J., Farmer, D., Packard, N. & Shaw, R. (1981) On determining the dimension of chaotic flows. *Physica* 3D:605-17. [AG]
- Garfinkel, A. (1983) A mathematics for physiology. *American Journal of Physiology* 245: (Regulatory, Integrative and Comparative Physiology) 14:R455-66. [arCAS]
- Gauld, A. & Shotton, J. (1977) *Human action and its psychological investigation*. Routledge & Kegan Paul. [RB]
- Gerstein, G. L. & Mandelbrot B. (1964) Random walk models for the spike activity of a single neuron. *Biophysical Journal* 4:41-68. [DHP]

References/Skarda & Freeman: Brain models chaos

- Gibson, J. J. (1979) *The ecological approach to visual perception*. Houghton Mifflin. [aCAS]
- Glass, L. & Mackey, M. (1979) Pathological conditions resulting from instabilities in physiological control systems. *Annals of New York Academy of Sciences* 316:214–35. [AG]
- Gonzalez-Estrada, M. T. & Freeman, W. J. (1980) Effects of carnosine on olfactory bulb EEG, evoked potential and DC potentials. *Brain Research* 202:373–86. [rCAS]
- Goodman, N. (1968) *Languages of art: An approach to a theory of symbols*. Bobbs-Merrill. [JAB]
- Grajski, K., Breiman, L., Viana Di Prisco, G. & Freeman, W. J. (in press) Classification of EEG spatial patterns with a tree-structured methodology: CART [Classification and regressive trees]. *IEEE Transactions in Biomedical Engineering*. [aCAS]
- Grassberger, P. (1986) Do climatic attractors exist? *Nature* 323:609–12. [AG]
- Grassberger, P. & Procaccia, I. (1983) Measuring the strangeness of strange attractors. *Physica* 9D:189–208. [aCAS]
- Gray, C. M. (1986) Centrifugal regulation of olfactory coding and response plasticity in the olfactory bulb of the conscious rabbit. Ph.D. thesis, Baylor University. [aCAS]
- Gray, C. M., Freeman, W. J. & Skinner, J. E. (1986) Chemical dependencies of learning in the rabbit olfactory bulb: Acquisition of the transient spatial-pattern change depends on norepinephrine. *Behavioral Neuroscience* 100:585–96. [aCAS]
- Grossberg, S. (1975) A neural model of attention, reinforcement, and discrimination learning. *International Review of Neurobiology* 18:263–327. [DSL]
- (1976) Adaptive pattern classification and universal recoding. II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics* 23:187–202. [rCAS]
- (1980) How does the brain build a cognitive code? *Psychological Review* 87:1–51. [aCAS, DSL]
- (1981) Adaptive resonance in development, perception, and cognition. In: *Mathematical psychology and psychophysiology*, ed. S. Grossberg. American Mathematical Society. [rCAS, SG]
- (1982a) *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Reidel. [SG]
- (1982b) Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychological Review* 89:529–72. [DSL]
- (1983) Neural substrates of binocular form perception: Filtering, matching, diffusion, and resonance. In: *Synergetics of the brain*, ed. E. Basar, H. Flohr, H. Haken & A. J. Mandell. Springer-Verlag. [SG]
- (1987a) *The adaptive brain. I: Cognition, learning, reinforcement, and rhythm*. Elsevier/North-Holland. [SG]
- (1987b) *The adaptive brain. II: Vision, speech, language, and motor control*. Elsevier/North-Holland. [SG]
- Grossberg, S. & Levine, D. S. (submitted) Neural dynamics of attentionally modulated Pavlovian conditioning: Blocking, interstimulus interval, and secondary reinforcement. [DSL]
- Grossberg, S. & Stone, G. O. (1986) Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review* 93:46–74. [SG]
- Guckenheimer, J. & Holmes, P. (1983). *Dynamical systems and bifurcations of vector fields*. Springer. [RT]
- Hebb, D. D. (1949) *The organization of behavior*. Wiley. [aCAS]
- Hinton, G. (1985) Learning in parallel networks. *Byte* 10:265. [aCAS]
- Hinton, G. & Anderson, J. A. (1981) *Parallel models of associative memory*. Erlbaum. [aCAS]
- Holden, A. V. ed. (1986) *Chaos*. Manchester University Press. [rCAS]
- Hopfield, J. 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* [aCAS, RR]
- Hopfield, J. J. & Tank, D. W. (1986) Computing with neural circuits: A model. *Science* 233:625–33. [aCAS, DSL, DHP]
- Julesz, B. (1984) A brief outline in the texton theory of human vision. *Trends in Neuroscience* 7:41–45. [aCAS]
- Kaczmarek, L. K. & Babloyantz, A. (1977) Spatiotemporal patterns in epileptic seizures. *Biological Cybernetics* 26:199. [AB]
- Kohonen, T. (1984) *Self-organization and associative memory*. Springer-Verlag. [aCAS, JAB]
- Lancet, D., Greer, C. A., Kauer, J. S. & Shepherd, G. M. (1982) Mapping of odor-related neuronal activity in the olfactory bulb by high-resolution 2-deoxyglucose autoradiography. *Proceedings of the National Academy of Sciences of the United States of America* 79:670–74. [aCAS]
- Lashley, K. S. (1942) The problem of cerebral organization in vision. In: *Biological symposia VII*, ed. J. Cattel. J. Cattel Press. [rCAS]
- Layne, S. P., Mayer-Kress, G. & Holzfuss, J. (1986) Problems associated with dimensional analysis of EEG data. In: *Dimensions and entropies in chaotic systems*, ed. G. Mayer-Kress. Springer. [AB]
- Levine, D. S. (1986) A neural network model of temporal order effects in classical conditioning. In: *Modelling of biomedical systems*, ed. J. Eisenfeld & M. Witten. Elsevier. [DSL]
- Lilly, J. C. & Cherry, R. B. (1951) Traveling waves of action and of recovery during responses and spontaneous activity in the cerebral cortex. *American Journal of Physiology* 167:806. [RMB]
- (1954) Surface movements of click responses from acoustical cerebral cortex of cat: Leading and trailing edges of a response figure. *Journal of Neurophysiology* 17:521–32. [RMB]
- (1955) Surface movements of figures in spontaneous activity of anesthetized cortex: Leading and trailing edges. *Journal of Neurophysiology* 18:18–32. [rCAS]
- Livanov, M. N. (1977) *Spatial organization of cerebral processes*. Wiley. [rCAS]
- Lorente de N6, R. (1934) Studies in the structure of the cerebral cortex. I: The area entorhinalis. *Journal von Psychologie und Neurologie* 45:381–438. [rCAS]
- Lorenz, E. (1963) Deterministic nonperiodic flow. *Journal of Atmospheric Sciences* 20:130–41. [AG]
- Mandell, A. J. (1986) Complexity versus disorder in the cardiac monitoring problem: A four-minute warning. Technical report, Mathematics Institute, University of Warwick. Coventry, England. [DHP]
- Marr, D. (1982) *Vision*. Freeman. [DCE]
- Marr, D. & Poggio, T. (1976) Cooperative computation of stereo disparity. *Science* 194:283–87. [DCE]
- Martinez, D. M. & Freeman, W. J. (1984) Periglomerular cell action on mitral cell in olfactory bulb shown by current source density analysis. *Brain Research* 308:223–33. [aCAS]
- Maturana, H. R. & Varela, F. J. (1980) *Boston studies in the philosophy of science*. Vol. 42: *Autopoiesis and cognition*, ed. R. S. Cohen & M. W. Wartofsky. Reidel. [GW]
- May, R. M. (1976) Simple mathematical models with very complicated dynamics. *Nature* 261:459–67. [DHP]
- McClelland, J. L. & Rumelhart, D. E. (1981) An interactive activation model of context effects in letter perception: Part I. *Psychological Review* 88:375–407. [JAB]
- McCulloch, W. S. & Pitts, W. H. (1943) A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5:115–33. [DHP]
- Misgeld, U., Deisz, R., Dodt, H. & Lux, H. (1986) The role of chloride transport in postsynaptic inhibition of hippocampal neurons. *Science* 232:1413–15. [aCAS]
- Moore, C. P., Perkel, D. H. & Segundo, J. P. (1966) Statistical analysis and functional interpretation of neuronal spike data. *Annual Review of Physiology* 28:493–522. [DHP]
- Moulton, D. G. (1976) Spatial patterning of response to odors in the peripheral olfactory system. *Physiological Reviews* 56:578–93. [aCAS]
- Mpitsos, G. J. & Cohan, C. S. (1986) Convergence in a distributed nervous system: Parallel processing and self-organization. *Journal of Neurobiology* 17:517–45. [DHP]
- Newberry, N. & Nicoll, R. (1985) Comparison of the action of baclofen with γ -aminobutyric acid on rat hippocampal pyramidal cells in vitro. *Journal of Physiology* 360:161–85. [aCAS]
- Nicolis, J. S. (1985a) *Hierarchical systems*. Springer. [AB]
- (1985b) *Chaotic dynamics of information processing with relevance to cognitive brain functions*. *Kybernetes* 14:167–73. [rCAS]
- Nicolis, J. S. & Tsuda, I. (1985) Chaotic dynamics of information processing: The “magic number seven plus-minus two” revisited. *Bulletin of Mathematical Biology* 47:343–65. [rCAS]
- Oono, Y. & Kohmoto, M. (1985) Discrete model of chemical turbulence. *Physical Review Letters* 55:2927–31. See especially the references to the work of Kuramoto, Yamada et al., most of which has been published in the journal *Progress in Theoretical Physics* over the past decade. [MAC]
- Pellionisz, A. & Llinás, R. (1979) Brain modeling by tensor network theory and computer simulation. The cerebellum: Distributed processor for predictive coordination. *Neuroscience* 4:323–48. [DHP]
- Perkel, D. H. & Bullock, T. H. (1968) Neural coding. *Neurosciences Research Program Bulletin* 6:221–348. [aCAS]
- Perkel, D. H., Gerstein, G. L. & Moore, G. P. (1967a) Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophysical Journal* 7:391–418. [DHP]
- (1967b) Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophysical Journal* 7:419–40. [DHP]

- Pylyshyn, Z. W. (1984) *Computation and cognition: Toward a foundation for cognitive science*. Bradford Books/MIT Press. [aCAS]
- Rapp, P. E., Zimmerman, I. D., Albano, A. M., Deguzman, G. C. & Greenbaum, N. N. (1985) Dynamics of spontaneous neural activity in the simian motor cortex. *Physics Letters* 110A:335. [AB]
- Rayleigh, Lord (John William Strutt) (1916) *Philosophical Magazine* 32:529–46. [RT]
- Riemann, B. (1902) *Gesammelte Werke, Nachträge*. Teubner. [RT]
- Roschke, J. & Basar, E. (in press) EEG is not a simple noise. Strange attractors in intercranial structures. In: *Dynamics of sensory and cognitive signal processing in the brain*, ed. E. Basar. Springer. [AB]
- Rosenblatt, F. (1962) *Principles of neurodynamics, perceptrons and the theory of brain mechanisms*. Spartan. [aCAS, DHP]
- Rössler, O. E. (1983) The chaotic hierarchy. *Zeitschrift für Naturforschung*. 38A:788–802. [arCAS]
- Ruelle, D. & Takens, F. (1971) On the nature of turbulence. *Communications in Mathematical Physics* 20:167 [RT]
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. (1986) Learning representations by back-propagating errors. *Nature* 323:533–36. [DHP]
- Rumelhart, D., McClelland, J. L. & PDP Research Group (1986) *Parallel distributed processing: Explorations in the microstructures of cognition*, vol. 1: *Foundations*. MIT/Bradford. [arCAS]
- Sampath, C. & Srinivasan, S. K. (1977) *Stochastic models for spike trains of single neurons*. Springer-Verlag. [DHP]
- Schuster, H. (1984) *Deterministic chaos: An introduction*. Physik-Verlag. [aCAS, RB]
- Searle, J. R. (1980) Minds, brains, and programs. *Behavioral and Brain Sciences* 3:417–57. [RB]
- Shaw, G. L., Silverman, D. J. & Pearson, J. C. (1985) Model of cortical organization embodying a basis for a theory of information processing and memory recall. *Proceedings of the National Academy of Sciences of the United States of America* 82:2364–68. [DHP]
- Shaw, R. (1984) *The Dripping Faucet as a Model Chaotic System*. Ariel Press. [rCAS]
- Sheer, D. (1976) Focused arousal and 40 Hz EEG. In: *The neuropsychology of learning disorders: Theoretical approaches*, ed. R. M. Knights & D. J. Baker, University Park Press. [rCAS]
- Shepherd, G. M. (1972) Synaptic organization of the mammalian olfactory bulb. *Physiological Review* 52:864–917. [aCAS]
- Sherrington, C. (1906) *The integrative action of the nervous system*. Yale University Press. [rCAS]
- (1940) *Man on his nature* (rev. ed. 1953). Cambridge University Press. [DHP]
- Skarda, C. A. (1986) Explaining behavior: Bringing the brain back in. *Inquiry* 29:187–202. [aCAS]
- Thom, R. (1981) Morphologie du sémiotique, recherches sémiotiques. *Semiotic Inquiry* 1:301–10. [RT]
- Turner, J. S., Roux, J. C., McCormick, W. D. & Swinney, H. L. (1981) Alternating periodic and chaotic regimes in a chemical reaction – experiment and theory. *Physics Letters* 85A:9–14 [AC]
- Viana Di Prisco, G. (1984) Hebb synaptic plasticity. *Progress in Neurobiology* 22:89–102. [aCAS]
- Viana Di Prisco, G. & Freeman, W. J. (1985) Odor-related bulbar EEG spatial pattern analysis during appetitive conditioning in rabbits. *Behavioral Neuroscience* 99:964–78. [aCAS]
- von Neumann, J. (1958) *The computer and the brain*. Yale University Press. [rCAS, DHP]
- Waddington, C. H. (1957) *The strategy of the genes*. Allen and Unwin. [RT]
- Walter, W. C. (1953) *The living brain*. Norton. [arCAS]
- Whitehead, A. N. (1960) *Process and reality*. Harper & Row. [rCAS, RT]
- Wiley, T. J. (1973) The ultrastructure of the cat olfactory bulb. *Journal of Comparative Neurology* 152:211–32. [aCAS]
- Wilson, D. A., Sullivan, R. M. & Leon, M. (1985) Odor familiarity alters mitral cell response in the olfactory bulb of neonatal rats. *Developmental Brain Research* 22:314–17. [rCAS, DHP]
- Winfree, A. T. (1980) *The geometry of biological time* (section 8D). Springer-Verlag. [MAC]
- Wolfram, S. (1984) Cellular automata as models of complexity. *Nature* 311:419–24. [DHP]
- Zucker, R. S. (1973) Changes in the statistics of transmitter release during facilitation. *Journal of Physiology (London)* 241:69–89. [DHP]