

EKG RESEARCH OF NEURAL DYNAMICS:
IMPLICATIONS FOR MODELS OF LEARNING AND MEMORY

Christine A. Skarda

CREA/Ecole Polytechnique
1, rue Descartes
75005 Paris, France

Walter J. Freeman

Department of Physiology-Anatomy
University of California
Berkeley, CA 94720

1. INTRODUCTION

What is learning and memory, or more specifically, at what level should neurophysiologists look to discover the neural processes used in the production of learned behavior and in controlling interactions with the environment? There are many changes that take place when a system learns or recalls what it has learned, but we suggest that neither learning or memory can be equated simply with change, and surely not every change at the neurophysiological level makes a difference from the system's point of view as far as its learned behavior is concerned. What neurophysiologists want to know is whether the "information" used by the system is to be found at the level of the single cell acting as an autonomous summing and switching device? Or should it be sought at the network level (Hebb 1949), at the metabolic level (Thatcher & John 1977) among masses of neurons, at the genetic level, or elsewhere?

The question is complicated at the outset in two ways. First, since the brain is a physicochemical machine operating simultaneously at multiple levels in a hierarchy, each change with learning constitutes a simultaneous modification at all levels of the hierarchy. Second, most of the changes are secondary and supportive or janitorial in nature. Our heuristic approach to this question is to assume at the outset that an optimal answer is to be found at the highest level of organization of the hierarchy, that which corresponds in its time and distance scales to the coordinate systems in which goal-directed behavior is elaborated by the species under study.

Many, and until recently most, neuroscientists have chosen the neuron as the element of function in learning and memory. This view has received its strongest support from physiological studies of the trains of action potentials elicited from single cells by structured stimuli. A cell that is driven maximally by a certain stimulus is said to be a feature detector for that stimulus, and its firing rate is said to be proportional to the probability that the feature being extracted or represented is present in the sensory field (Barlow 1972). The colloquial and somewhat pejorative title of the "grandmother cell hypothesis" has been used to refer to the cell or "college of cardinal cells" that fire when one sees or thinks of one's grandmother. Such cells that serve as the repository for cognitive information at the atomic level of a word are proposed to exist as the basis for memory in each sensory and associative system, and to take the form of "command cells" in motor and premotor systems. The resulting task for an explanation of brain function has been to determine the patterns of forward connections among cortical neurons

that can explain how complex features such as line segments and angles are synthesized from simple features such as on-center off-surround cells (Hubel & Wiesel 1962).

Working downward, some researchers have suggested that the proper level for understanding lasting changes with learning is at the neuronal organelle, such as the synapse (Hebb 1949) and the dendritic spine (Rall & Rinzel 1973; Lynch 1986). Kandel and Schwartz (1982) designate the best level to be the biochemical operations in the synapse, particularly the modifications of calcium flux and the effects on protein assembly and deployment in the synaptic complex. Kandel proposes to extend the search into the nucleus of the neuron on the premiss that any permanent change in the protein-making capability of a neuron that subserves learning may entail a change in the genome, or at least a modification of the RNA in the Nissl substance. This brings us full circle back to the hypothesis of Hyden (1960) that fostered a cottage industry of the sixties in which high school students across the country sought to enrich the lives of flatworms by feeding them the RNA of their educated brethren.

Working in the opposite direction, researchers have postulated that the changes with learning, although constituted by or based on cellular and molecular modifications, are spatially widely distributed and must be understood first at the level of neuronal networks (Amari 1983; Grossberg 1981; Hérault & Jutten 1988; Hinton 1985; Hopfield & Tank 1986; Rumelhart et al. 1986). For example, if in accordance with the Hebb postulate there is strengthening of some synapses and decay of others in an interactive population of cells leading to the formation of a nerve cell assembly (NCA), then in some cells there must be stronger synaptic actions. These will require more energy in the form of ATP so that more mitochondria must be made with the result that all mitochondrial proteins will be increased in the modified tissue. If the numbers of synapses are increased by cloning, or if their size or the length of spines is increased, then more membranal molecules are required. The metabolic demands may increase the activities of glial cells which have more work to do keeping the neural electrolytes in balance. These changes have led some to posit that learning and memory take place at the metabolic level among huge masses of neurons in the brain (Thatcher & John 1977). We suggest, however, that although these and many other changes support learning, they do not explain it or constitute the essence of it.

The issue is further complicated by the circumstances by which learning is induced in normal subjects. Some form of repeated stimulation is required that in itself imposes metabolic and electrolytic demands on neurons and glia. Concomitantly, there are contextual and background processes no less well defined at present than those subserving associative learning. The unconditional stimulus that serves for reinforcement must have widespread effects of diverse kinds. Typically, both conditioned and unconditioned stimuli are made rather strong so as to attract the attention of a subject; thus, they impose demands for range compression and normalization, such as occurs by light adaptation in the visual system and logarithmic conversion of stimulus intensity in the olfactory system (Freeman 1975). The fact that these are journeyman and janitorial processes of a routine and nonproblematic kind does not sweep them out of the way when researchers are probing for the structural and chemical concomitants of learning. More likely these are the first to be stumbled over.

How, then, can we hope to ferret out and critically identify the neural changes with learning that signify and label the processes of associative memory? We propose that the crucial step consists in observing and measuring the neural activity patterns of sensory systems before and after a subject has learned to discriminate two or more sensory stimuli, to identify the precise nature of the differences in activity patterns that serve to distinguish and properly classify the neural events in respect to the discriminanda, and to construct a

dynamic model that suffices to replicate the observed pattern differences. This dynamic model is crucial, in our opinion, because it serves to specify the type of synapse that undergoes modification with learning, where it is located, when it changes, and the manner in which the synaptic change subserves the global network change leading to altered goal-directed behavior, particularly how each neural activity pattern, that supports the appropriate form of behavior, is selected and brought to fruition in the presence of the discriminated stimulus.

We believe that the neural network approach to these issues is valuable, incorporating as it does earlier insights on synaptic changes associated with learning and memory, but it is still an open question as to how the mechanisms of synaptic change and the formation of neural networks function to produce learned behaviors. Is the formation of patterned activity in a network via the underlying processes of synaptic change sufficient to explain learning and to produce a memory and/or memory recall in a system? Our data suggests that more is involved. In what follows we describe the neural dynamics responsible for learning and memory, some respects in which our model differs from recent "neural net" models, and some implications of our findings for further research.




FIGURE 1

EEG waves are shown from a selected channel in the olfactory bulb of a waking rabbit (top trace) and from the prepyriform cortex to which the olfactory bulb transmits its output. Receptors in the nose sending their axons to the bulb are activated by air inflow during each inhalation (upward peaks of middle trace). The surge of input on many axons causes this "respiratory wave" that manifests excitation of the bulb. Near the crest of excitation the bulb breaks into its own oscillation giving rise to a burst riding on the respiratory wave. Both the wave and the burst are transmitted to the prepyriform cortex where they are mixed with other activity that is transmitted to other parts of the brain. The burst has the same frequency everywhere in the bulb; the modulation of its amplitude in the spatial dimensions of the bulb provides a pattern that contains and transmits the odor-specific information subserving discriminative responding and the appropriate behavior, in this case a sniff that is induced by a warning odor paired with shock in a classical aversive conditioning paradigm. (From Freeman & Schneider 1982)

2. NEURAL DYNAMICS MANIFEST IN EEGS

The model of learning and memory presented here concerns paleocortical dynamics; more specifically it concerns preattentive sensory processing in the olfactory bulb. It is based on experiments in which thirsty rabbits were conditioned to lick (CR+) in response to an odorant (CS+) followed after 2 s by delivery of water (UCS), and just to sniff (CR-) in response to an unreinforced odorant (CS-). The model is derived from studies of changes in the waveform of evoked potentials in the olfactory system of rabbits recorded with a chronically implanted 8 x 8 array of electrodes (spacing: 0.5 mm) covering approximately 20% of the surface of the olfactory bulb, and on replication of these waveforms by nonlinear differential equations simulating the dynamics of the olfactory bulb, anterior olfactory nucleus and prepyriform cortex. The typical pattern of bulbar electroencephalogram (EEG) trace (Figure 1) was a slow wave common to all 64 channels on the crest of a respiratory wave with a burst of oscillation in the gamma range (35-90 Hz). Spatial analysis of the EEG revealed that odor specific information exists in spatial patterns of amplitude of this oscillatory burst.

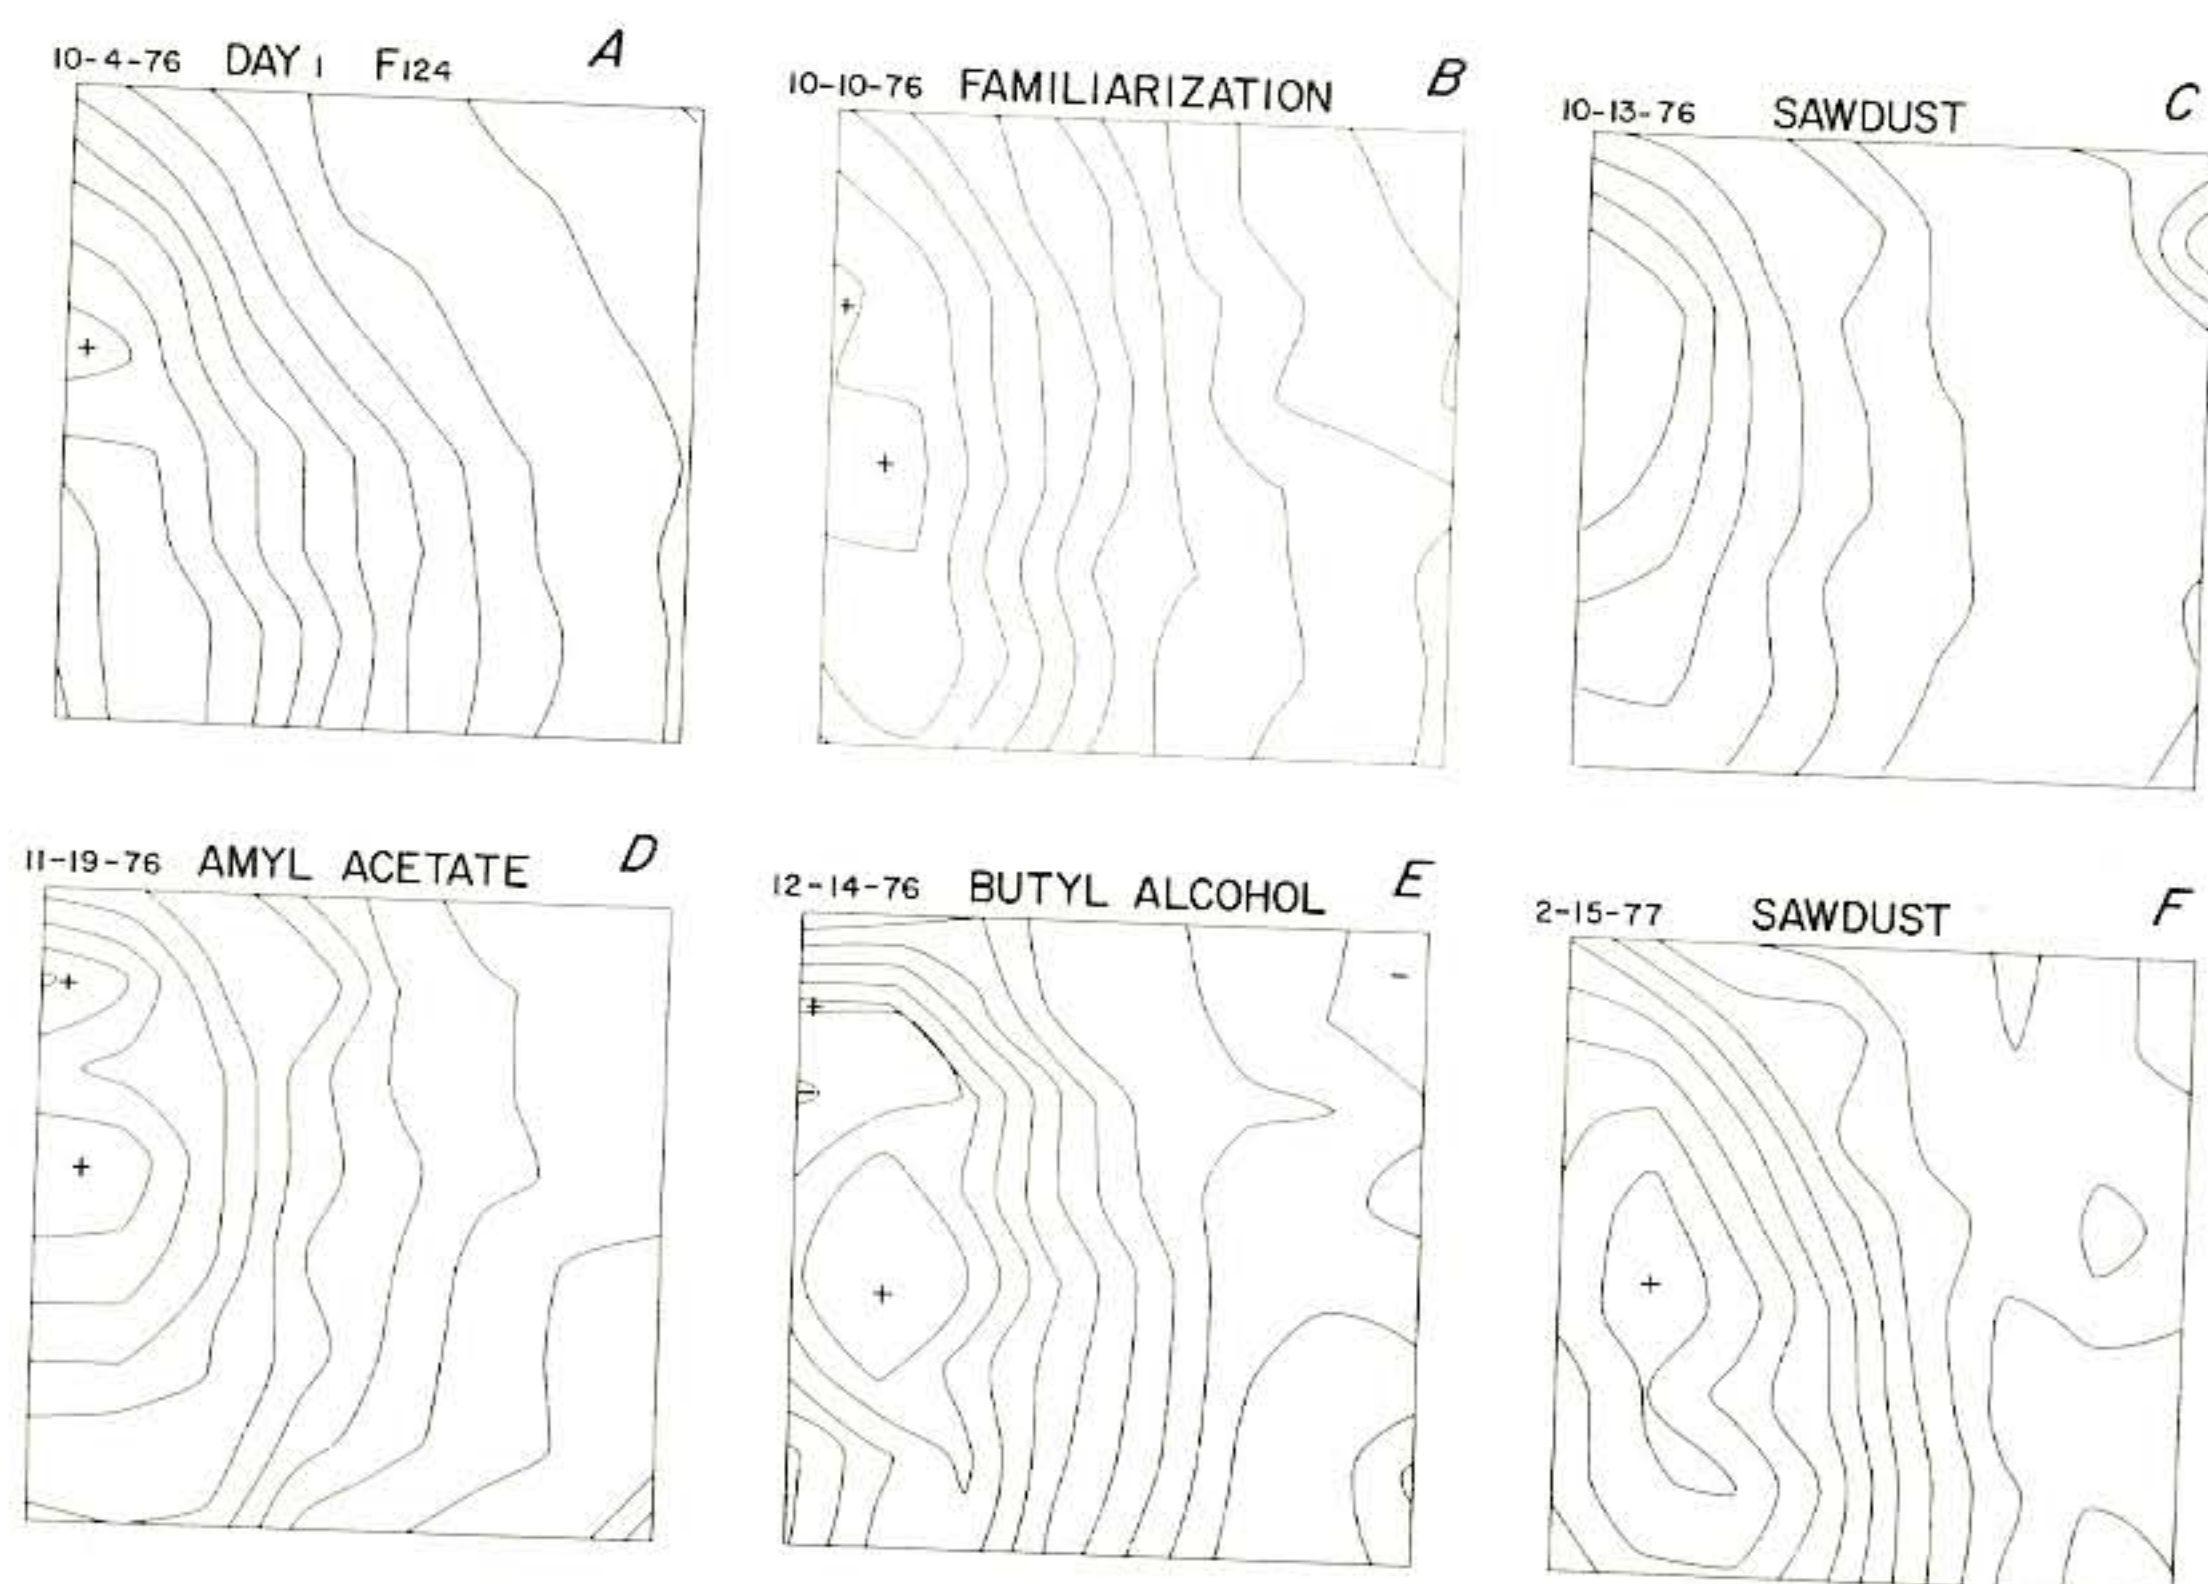


FIGURE 2

Each burst has the same wave form (frequency) over the entire bulb or cortex that differs in amplitude from place to place. Hence the odor-specific information can be expressed by the spatial pattern of amplitude. These examples from a rabbit olfactory bulb show the patterns by means of contour plots from an array of 64 EEG bursts recorded from an 8 x 8 array at different stages of conditioning. At top left is the pattern first seen after recovery from surgery. Next is the pattern after familiarization with the recording and test apparatus. Then follows the pattern associated with the first test odor (sawdust), followed by two other odorants, each sequentially paired with shock. The last (lower left) frame shows that return to the original odorant (sawdust) leads to a new pattern and not recurrence of the original one (upper right). This memory is not tied to our catalog of chemicals. It is an on-going experiential process, fully associational, that has a history that cannot exactly repeat itself. (From Freeman & Schneider 1982)

Analysis of the EEG traces revealed that in the background state (before conditioning) every trace had the same temporal waveform, and that the amplitude differed between channels to form a relatively constant spatial pattern that was easily identified with a particular animal. The patterned activity of this background state remained constant until odorant conditioning was undertaken (Figure 2). New patterns emerged only with reinforced odorants and remained stable within and across sessions provided S-R contingencies were not changed. Multiple patterns emerged under discrimination conditioning. These patterns were globally distributed; it was not possible to localize the information in the stable spatial patterns that served to classify events correctly to a particular subset of channels. The task of modeling was to produce a simulation of the robust background state and the emergence of globally distributed, odor specific spatial patterns within constraints set by known characteristics of olfactory functioning.

How does learning occur in the olfactory bulb? The model based on these studies can be briefly characterized as follows: Excitatory neurons, synaptically linked by bidirectional synapses, are co-activated pairwise by a CS+ strengthening their joint synapses in accordance with Hebb's rule (1949). A nerve cell assembly (NCA) is thus created by strengthened excitatory synapses comprising perhaps 1-5% of the total in the olfactory bulb. Thereafter, excitation of any portion of the NCA by receptors sensitive to a particular odor tends to activate the whole assembly. Formation of the NCA under conditioning is the essential part of the lasting physiological changes associated with learning in the olfactory system, but more is involved in odor recognition and in the process of retrieval or memory.

Physiological studies reveal that during inhalation input to the bulb augments exponentially the strength of interaction of the entire bulb, not just a selected subset of neurons. The model indicates that during late inhalation, bulbar neurons are pushed far from their initial low energy state resulting in a state change or bifurcation. It is hypothesized that the NCA operates at the moment of choice when the surge of receptor input strongly forces the bulb away from its rest state to a new activity pattern. The hypothesized role of the synaptic changes leading to the formation of the NCA is threefold: (1) to amplify and stereotype the small input received on any given inhalation, (2) to produce a locally disseminated but low density activity pattern from its interaction with the stimulus, and (3) to provide the crucial mechanism for mediating the emergence of the globally distributed odor-specific activity pattern at the time of state change. The NCA is thought to determine the "basin of attraction" for the system leading to the emergence of a globally distributed, odor specific activity pattern that is mathematically expressed as a limit cycle attractor. It is only with this state change that 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 stereotypical activity pattern that is mediated by the NCA but goes far beyond it. It is hypothesized that this globally distributed, stereotypical pattern constitutes odor recognition and memory for the rest of the system; this is the pattern that is made available to the rest of the brain, and that is behaviorally relevant for the kind of correlations that are usually associated with learning and memory.

EEG data and the resulting model reveal that the olfactory bulb operates in number of states. During late exhalation and early inhalation intrinsic interaction strength is low, and the activity of afferent neurons is imposed on bulbar neurons that can accept information and maintain it by local firing. The model indicates that this background state is a low level chaotic state in which is embedded the locally disseminated activity pattern of the NCA. On bifurcation the bulb converts to a transmitting mode in which internal interaction is high and bulbar neurons no longer respond to receptor input. In this state, information carried by each neuron is disseminated over the entire

bulb, is integrated by every neuron in the bulb, and is sent out of the bulb to the cortex which, in turn, sends feedback to the bulb. Upon exhalation the bulb returns to its low level chaotic state.

We believe that our research and the resulting model are significant for researchers interested in investigating the neurophysiological basis of learning and memory for at least two reasons. First and foremost, the model suggests in outline the way in which synaptic change and the NCA fit into the larger framework of memory formation and recall. We hypothesize that the formation of a NCA and its resulting locally distributed activity pattern is necessary, but not sufficient, to explain memory, and that research focused exclusively on this process will miss the role played by the globally distributed, self-organized, internally generated activity pattern that in the olfactory system constitutes memory but that does not itself require further synaptic alteration. Second, this "bigger picture" allows us to distinguish from the point of view of system dynamics those processes primarily concerned with learning from those that subserve memory. We hypothesize that the synaptic changes occurring under reinforcement leading to the formation of a NCA that is locally disseminated, are the essential long-term material changes for learning in the olfactory system, while globally distributed activity patterns taking place across the entire bulb are the hallmark of memory formation and the recall process.

3. FUNCTIONAL ARCHITECTURE AND NEURAL NET MODELS

The model of sensory functioning in the olfactory system as a distributed, self-organizing process that is mathematically expressible in terms of nonlinear dynamics resembles some of the neural net models developed in recent connectionist research, namely that class of connectionist systems characterized by self-directed dynamical models (Amari 1983; Freeman 1987; Grossberg 1981; Hopfield & Tank 1986; Kohonen 1984). In such models learning is a distributed process that takes place without a teacher by gradually varying connection strengths among units comprising the network. Learning is, thus, in the connections and knowledge or memory is in the capacity of the network to arrive at similar input-output pairings. But in such systems learning is not only globally distributed, it is self-organized. Unlike so-called PDP systems (Rumelhart et al. 1986) that rely on feed-forward connectivity and back propagation for error correction and that have their "goals" externally imposed by the system's operator in the form of a "teacher" or set of correct answers with reference to which the output of the system is corrected via error correction, self-organizing systems do not use internal representations of any sort to generate a more ordered state leading to reliable forms of interaction with the environment from less ordered initial conditions. Thus, these models differ significantly from models based on conventional computer technology: not only do self-organized connectionist systems use distributed, highly parallel processing, they do not compute with symbolic tokens, they do not use "teachers" to specify the desired goal, and they do not calculate according to externally specified ad hoc rules.

Even with self-organizing systems, however, there is not a complete isomorphism between actual neural dynamics and the so-called neural net models of connectionist research. Some of the differences have been reviewed in a recent paper (Skarda & Freeman 1987), but there are several features of the model of olfactory dynamics that are of special interest in the present context.

(1) Neural net models of learning and memory focus exclusively on the mechanisms of synaptic change to the exclusion of other possible learning mechanisms. In this respect they reflect important insights into neuronal functioning gained from single unit research. However, while the brain clearly uses the mechanisms of synaptic change (both transient and long term anatomical

ones), it seems safe to assume that a system that is extremely economical in many respects would not overuse this highly valuable resource. One suggestion is that future neural net modelers might be wise to try to use changes based on the Hebb rule sparingly, combining it with other types of learning mechanisms.

In the olfactory model discussed above synaptic change is only part of the story, although admittedly an essential part. Learning results in synaptic changes that lead to the formation of a NCA. But it is the resulting patterned activity of the NCA rather than further synaptic changes that mediates the recall of a global, self-organized state in the whole bulb, a state that does not itself require further Hebb-type synaptic changes. Finally, this global odor-specific activity pattern is hypothesized to play a role in determining activity at later stages of information processing (Freeman & Skarda 1985).

(2) This raises a related issue, the role of hierarchically arranged neural processing. Neural net models of the Hopfield and Hinton varieties tend to be 'simple' in the sense that they involve no real architectural complexity. Recently, however, connectionist modelers have begun to think in terms of creating architecture that will enable their models to handle problems that traditional computational systems were able to handle, e.g., inference and sequential processing, but that neural net models have not yet tackled. The neural model based on EEG data encourages this trend and indicates that in the olfactory system a hierarchically arranged, self-organized neural dynamics is responsible for odor recognition.

In the olfactory system architecture is postulated to play a crucial role both within the olfactory bulb itself and among individual subsystems. While the limbic system as a whole may be viewed as a hierarchy of interacting, self-organized states, within the olfactory bulb itself there exists a hierarchy of self-organized stable states ranging from deep anesthesia to seizure states (Figure 3). Depending on various neural conditions, and in the model on parameter changes, the system behaves differently even though the basic neural dynamics (and equations) are the same in all states. In the model the (chaotic) background state can give way to a low energy state that can be mathematically represented as a point attractor for states of deep anesthesia or death, or to a limit cycle attractor for each learned odor, or to a high energy chaotic attractor in the seizure state (Figure 4). It is postulated that in a motivated animal all of the learned tendencies to produce patterned behavior associated with particular odors would be available with each inhalation thanks to the chaotic nature of the background state that gives access to the entire state space in a controlled manner, and with inhalation a selection is made based on activity in a NCA or its absence. The availability of these odor specific patterns (attractors) vanishes with exhalation and the system is freed to learn a new odor or to recognize one that is already learned.

We have postulated that segregation into a hierarchy may be necessary for several reasons (Freeman & Skarda 1985). Included here is evidence from anatomical studies on patterns of interaction and interconnection in the limbic system that reveal that successive levels of organization and integration are involved in the formation and operation of this system. In addition, the orderly sequence of transitions between stable states that we have observed requires that the details of dynamic activity in each subsystem or part be collapsed into a small number of variables and parameters, such that a state variable at a higher level can serve as a control parameter at lower levels.

(3) There is a final difference between current neural net models and neural dynamics in the olfactory system worth noting in the present context. Most neural net models are not "selective", i.e., they process whatever information is received by receptor level neurons. When presented with a pattern neural net models begin to process the resulting receptor level pattern. This is not the case in the olfactory system. In the bulb, receptor level activity only leads to the formation of a NCA under reinforcement (Freeman 1981). In unmotivated animals receptor level activity still occurs,

but no corresponding changes are recorded in the olfactory bulb and no globally distributed patterned activity emerges that signals odor recognition and recall to the rest of the brain. This difference suggests that actual neural dynamics are not reactive, like connectionist models. We suggest that they be viewed as 'active', in the sense that they incorporate an internal neural process that determines whether receptor level activity will have any effect on the rest of the system. Learning, on this model, is not simply reacting to receptor input that triggers changes in internal neural dynamics, it is selective or discriminative from the outset.

4. IMPLICATIONS FOR LEARNING AND MEMORY RESEARCH

In the olfactory system, as the model suggests, information processing involves single units, cell assemblies, and mass action among large populations located in regions or neural subsystems. At the receptor level single unit activity responding selectively to odorants plays an important role in establishing mutual excitatory activity among pairs of neurons in the bulb leading to the formation of the NCA; the formation of the the NCA, in turn, is the essential physiological correlate of learning involving locally distributed synaptic




FIGURE 3

Four EEG traces are shown from a representative channel of an array in a rat that serve to identify four behavioral states and the accompanying dynamic states of the olfactory system. The lowest trace is from a state of deep anesthesia induced by a surgical dose of barbiturate. The next is from a waking animal that is motionless and unmotivated. This is a state of low level chaos, distinct from the equilibrium or unchanging state under deep anesthesia. The third state is the "wave and burst" pattern shown above the concomitant record of respiration. The top trace shows a seizure state in which normal behavior is in abeyance, the animal goes into a trance, and a spike train occurs a 3/second. (From Freeman 1987)

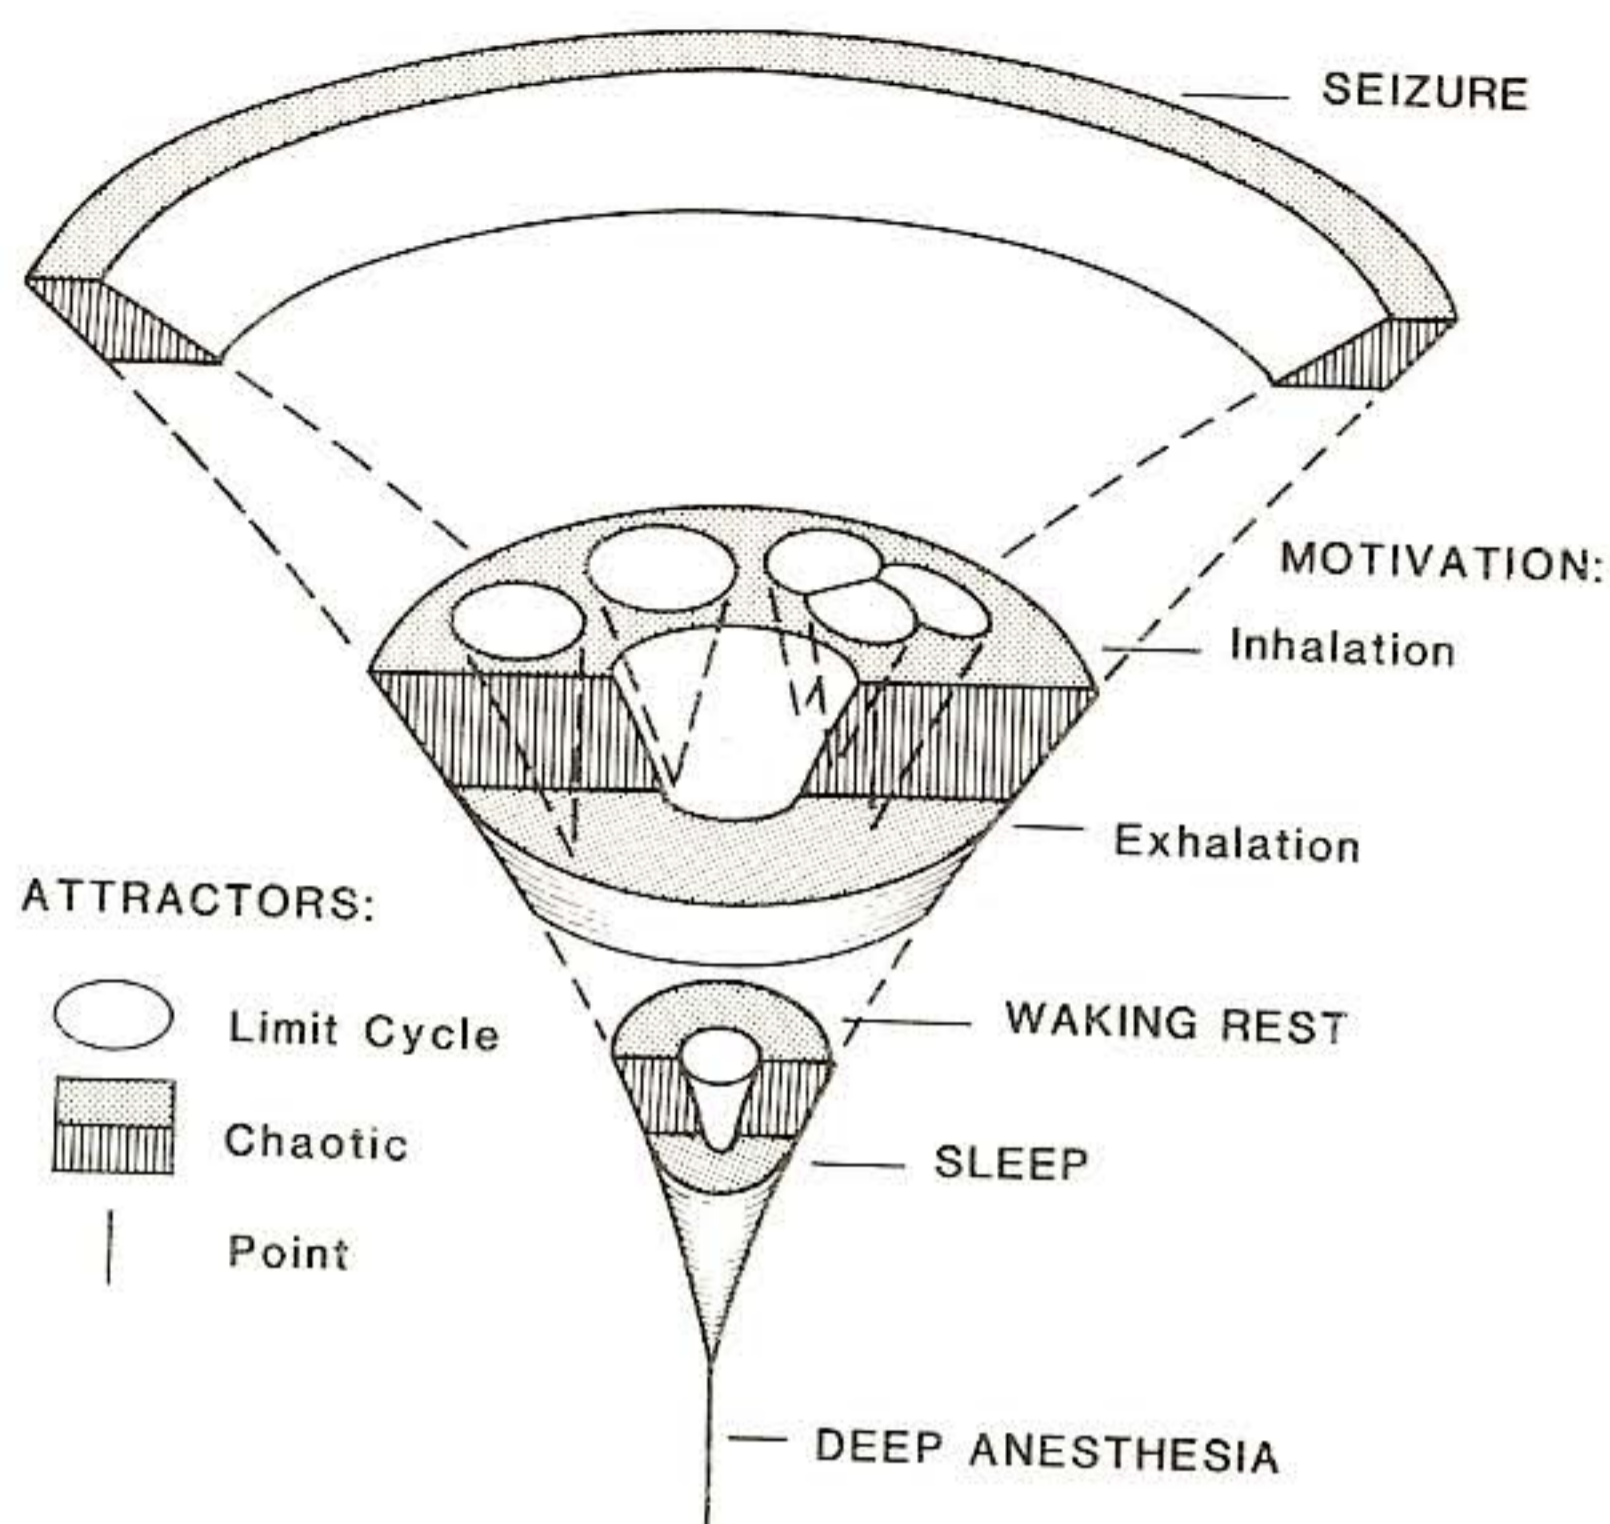


FIGURE 4

The mathematical model that allows us to simulate the EEG patterns shown in Figure 3 suggests that a hierarchical set of states exists for the olfactory bulb, anterior nucleus and prepyriform cortex, ranging from equilibrium at the lowest energy state up through low level chaos by the so-called Ruelle-Takens route. Under motivational activation the system goes to a higher energy state of chaos, from which it can be forced to bifurcate into one of a set of limit cycle attractors that have been formed by discriminative conditioning. These attractors are latent and come to exist only with inhalation; they vanish upon exhalation. With yet further energization, provided by intense electrical stimulation, the system enters a seizure state that is also chaotic. (From Freeman 1987)

changes and ultimately global patterned activity; finally, discriminatory behavior and memory for a particular odor depend upon the existence of patterned output from the olfactory bulb which is the product of globally distributed, self-organized behavior among all the neurons in the bulb.

This implies that as far as experimental approaches are concerned various techniques will be required to understand the complete story of the neurophysiological basis of learning and memory. For example, study of single unit properties alone cannot address the coordinated and distributed changes in large populations of neurons or in the NCAs that are hypothesized to underlie memory processes. On the other hand, EEG studies cannot describe the changes leading to the formation of a NCA. We believe it is essential for researchers in the field to recognize the level at which their investigations are pitched and to understand how the phenomena they study fit into the "bigger picture" of neural processing underlying learning and memory.

We are convinced that having a "big picture" of the neural dynamics responsible for learning and memory is essential for researchers interested in answering the crucial question as to where to look for the information gained when an animal learns or remembers something. We, for example, have found our model of

olfactory system dynamics useful in rejecting two recent views of the neurophysiological basis of memory. One says that memory traces are stored in particular neurons or in neural circuits involving identifiable neurons (Barlow 1972), while the other view claims that each memory involves an enormous system of cooperating neurons, perhaps as much as one tenth of all of the neurons in the brain (John et al. 1986). Our experimental data and the resulting model argue against both views. On the one hand, the model indicates that what is 'stored' in the brain is perhaps best understood as a tendency to engage in a particular form of patterned, self-organized activity, and that this involves neurons locally linked synaptically into NCAs, but also large globally distributed populations of neurons. Yet our evidence also indicates that the process of learning and memory in the olfactory bulb is not infinitely distributed, even though it may involve metabolic changes throughout the brain. The model suggests that there is a difference between neurons actually involved in learning and memory processes, and a larger population of neurons in the nervous system which receive input from neurons activated when the system learns or remembers something. This larger population is affected by activity directly responsible for memory and recall as the consequences of discrimination are realized, but the information relevant to the discrimination is not contained in these widely distributed metabolic changes.

5. CONCLUSIONS

In the olfactory system odor learning and recall involve a hierarchically arranged set of processes that are internally generated and cannot be understood as simple responses to external stimuli. We say that they are "self-organized". Our data indicate that learning occurs only under reinforcement in motivated animals, and that it involves chemically mediated synaptic change leading to the formation of a NCA sensitive to a particular odor (Gray et al. 1986). Once neurons have had their synaptic interconnections strengthened by this learning process they tend to preferentially excite one another, such that excitation of a subset of the neurons comprising the NCA leads to a stereotypical pattern of activity in the whole neural network. Some of the recent connectionist or neural net models have attempted to describe the formation of such networks of excitatory neurons (e.g. Hopfield and Tank 1986) but not with locally endogenous oscillatory activity. In this respect, actual neural dynamics differ from current connectionist models.

Evidence from EEG studies of olfactory functioning indicates that odor recognition involves more than activity in a NCA. Input to the bulb secondarily excites all bulbar neurons after excitation of the subset of neurons involved in the NCA. When a critical threshold is reached a state change occurs via a self-organized process in which the entire bulb changes to a globally distributed stereotypical pattern of oscillatory activity. Activity in the locally distributed NCA mediates the choice of this globally distributed patterned activity, but further synaptic change is not required once the NCA is established. We hypothesize that evocation of this globally distributed, oscillatory pattern constitutes odor recognition and memory for the system, since this is the pattern that is made available to the rest of the brain.

Analysis of EEG traces yields important insights into neural structure and function (see also Babloyantz & Destexhe 1987), insights that cannot be discovered using other methods. In self-organized systems, global dynamics cannot be explained or predicted from studies of the parts that constitute the system (Crutchfield et al., 1987). As a result, single unit studies will not be able to explain learning and memory in the brain, although they can investigate the mechanisms involved in synaptic change leading to the formation of a NCA. We suggest that an adequate explanation of learning and memory awaits cooperative involvement of researchers using experimental methods pitched at a higher or lower level of organization than that of the neuron.

REFERENCES

- Amari, S. (1983) Field theory of self-organizing neural nets. *IEEE Transactions on Systems, Man, and Cybernetics*. SMC-13, No. 5: 741-48.
- Babloyantz, A. & Destexhe, A. (1987) Chaos in neural networks. *Proceedings of First International Conference on Neural Networks*. San Diego, CA. (unpublished).
- Barlow, H. (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1: 371-94.
- Crutchfield, J. et al (1987) Chaos. *Scientific American*, 255: 46-57.
- Freeman, W. (1975) *Mass Action in the Nervous System*. Academic Press.
- Freeman, W. (1981) A physiological hypothesis of perception. *Perspective in Biology and Medicine*, 24: 561-592.
- Freeman, W. (1987) Analytic techniques used in the search for the physiological basis of the EEG. In *Handbook of Electroencephalography and Clinical Neurophysiology*, ed. A. Gevins & A. Remond. Elsevier.
- Freeman, W. & Schneider, W. (1982) Changes in spatial patterns of rabbit olfactory EEG with conditioning to odors. *Psychophysiology*, 19: 44-56.
- Freeman, W. & Skarda, C. (1985) Spatial EEG patterns, nonlinear dynamics and perception: The neo-Sherringtonian view. *Brain Research Reviews*, 10: 147-75.
- Gray, C., Freeman, W. & Skinner, J. (1986) Chemical dependencies of learning in the rabbit olfactory bulb: Acquisition of the transient spatial pattern depends on norepinephrine. *Behavioral Neuroscience*, 100(4): 585-596.
- Grossberg, S. (1981) Adaptive resonance in development, perception, and cognition. In *Mathematical Psychology and Psychophysiology*, ed. S. Grossberg. American Mathematical Society.
- Hebb, D. (1949) *The Organization of Behavior*. Wiley.
- Herault, J. & Jutten, C. (1988) Independent components analysis in neural networks. This volume.
- Hinton, G. (1985) Learning in parallel networks. *Byte*, 10: 265.
- Hopfield, J. & Tank, D. (1986) Computing with neural circuits: A model. *Science*, 233: 625-33.
- Hubel, D. & Wiesel, T. (1962) Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160: 106-154.
- Hyden, H. (1967) ed., *The Neuron*. Elsevier.
- John, E.R., Tang, Y., Brille, A., Young, R., & Ono, K. (1986) Double-labeled metabolic maps of memory. *Science*, 233: 1167-1175.
- Kandel, E.R. & Schwartz, J.H. (1982) Molecular biology of learning: Modulation of transmitter release. *Science*, 218: 433-443.

Kohonen, T. (1984) *Self-Organization and Associative Memory*. Springer-Verlag.

Lynch, G. (1986) *Synapses, Circuits, and the Beginnings of Memory*. MIT Press.

Rall, W. & Rinzel, J. (1973) Branch input resistance and steady attenuation for input to one branch of a dendritic neuron model. *Biophysics Journal*, 13: 648-688.

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

Skarda, C. & Freeman, W. (1987) How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences*, 10: 161-173.

Thatcher, R. & John, E.R. (1977) *Foundations of Cognitive Processes*. Erlbaum Associates.