

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS

The copyright law of the United States [Title 17, United States Code] governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that use may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgement, fulfillment of the order would involve violation of copyright law. No further reproduction and distribution of this copy is permitted by transmission or any other means.

Bloomsburg University ILL



ILLiad TN: 137068

Borrower: PMC

Lending String: *PBB,DXU,EIB,HVC,LAF

Patron: Loring, Patricia

Journal Title: The Behavioral and brain sciences.

Volume: 10 **Issue:** 2

Month/Year: 1987**Pages:** 181-183

Article Author: Ronald Rosenfeld, David S.
Touretzky and the Boltzmann Research Group

Article Title: Connectionist Models as Neural
Abstractions; Commentary on "Brains Make Chaos
to Make Sense of the World" by C.A. Skarda and
W.J. Freeman

Imprint: [New York] Cambridge University Press

ILL Number: 107099875



Call #: bounds

Location: 4th

ODYSSEY

Charge

Maxcost: 35.00IFM

Shipping Address:

Carnegie Mellon University - ILL
Eng and Sci Library
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
IDS 187A01

Fax: (412)681-1998

Ariel: 128.2.114.135

Email: es82@andrew.cmu.edu

Odyssey: 128.2.20.146

Bloomsburg University of Pennsylvania

Interlibrary Loan

NOTICE OF COPYRIGHT RESTRICTIONS

The copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyrighted materials.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be "used for any purpose other than private study, scholarship, or research". If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that user may be liable for copyright infringement.

This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve violation of copyright law.

No further reproduction or distribution of this copy is permitted by electronic transmission or other means.

Problem report

If you have a problem with the delivery of the requested items, please contact us with the following information:

Missing pages: _____

Edges cut off: ____ Left ____ Right ____ Top ____ Bottom

Unable to read: ____ Blurry ____ Too dark ____ Too Light

Other: _____

Contact us:

Phone: 570-389-4218

Ariel: 148.137.188.215

Fax: 570-389-3895

Odyssey: 148.137.188.233

Email: ill@bloomu.edu

OCLC: PBB

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 (Mptsos & 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 specific

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 (I) 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 "sailance" (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 longstanding 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