

LA NUOVA CRITICA

NUOVA SERIE

9-10

LA COSTRUZIONE DELLA NATURA

3°

Organizzazione, complessità e modelli cognitivi

FRANZ M. WUKETITS, Cognition: a non-adaptationist view.

GIUSEPPE DEL RE, Struttura e organizzazione nelle scienze della complessità.

ELENA GAGLIASSO, Forma, struttura, informazione: annotazioni tra teoria e storia.

CHRISTINE A. SKARDA, Understanding perception: selforganizing neural dynamics.

ARTURO CARSETTI, Teoria della complessità e modelli della conoscenza.

CHRISTINE A. SKARDA

UNDERSTANDING PERCEPTION:
SELF-ORGANIZING NEURAL DYNAMICS

1. Introduction

Perception is mischaracterized when described as a neural reaction to receptor input or to «what is out there». One class of objections to this stimulus-response model of physiological functioning appeals to a nonmaterial element or «ghost» in the machine. This dualistic view does not make physiological sense and is unnecessary; there is no need to introduce a nonbiological element to explain perceptual processing.

Perception is not produced by an internal ghost, but neither is it adequately characterized by the currently popular alternative to dualism and to stimulus-response theories, representationalism. According to representationalist accounts, perception is a physiological process whose goal is to create an internal model or representation of the environment. I shall argue that once the self-organizing mechanisms responsible for perceptual experience are understood, representationalism, too, can be abandoned as a mischaracterization.

Yet a positive characterization of perceptual processing in the brain requires more than the simple recognition that self-organization plays a role. Recent connectionist research demonstrates that there are a number of ways to incorporate self-organizing dynamics in an account of perceptual processing. Review of recent data on neural dynamics shows that brains use self-organization in a way that no current connectionist model does, and that there are important implications here for how we view perception. In what follows I will briefly review connectionist research and recent findings gathered from electroencephalogram (EEG) research on perceptual processing in the

olfactory system. My aim is to replace traditional accounts of perceptual processing with a new picture of perception.

2. *How the Brain Perceives*

The theory of brain function presented here is based on EEG studies of preattentive sensory processing in the olfactory bulb. This data has been discussed elsewhere in detail (Freeman & Skarda 1985), and can be summarized briefly here. The model concerns paleocortical dynamics and is based on experiments in which thirsty rabbits were conditioned to lick (Cr +) in response to an odorant (CS +) followed after 2 seconds by delivery of water (UCS), and just to sniff (CR-) in response to an unreinforced odorant (CS-). The findings are derived from the study of changes in the waveform, of evoked potentials recorded using a chronically implanted 8 x 8 array of electrodes 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. It was found that the typical bulbar EEG trace is a slow wave common to each of the 64 channels on the crest of a respiratory wave with a burst of oscillation in the gamma range (35-90 Hz).

Analysis revealed that it is possible to correlate the spatial patterns of amplitude of this oscillatory burst with specific odors. In the background state (before conditioning) every trace has the same temporal waveform, but the amplitude differs between channels to form a relatively constant spatial pattern that can reliably be identified with a particular animal. The patterned activity of this background state remains constant until odorant conditioning is undertaken, at which point new patterns begin to form. These new patterns emerge, however, only with reinforced odorants and remain stable within and across sessions only when the reinforcement contingencies are not changed; the significance of this will become clear below. Under discrimination conditioning multiple patterns emerged. These patterns were globally distributed, i.e., it was not possible to localize the information that served to classify events correctly in the stable spatial patterns to a particular channel or subset of channels. The aim of modeling was to simulate the background state and the emergence of global-

ly distributed, odor specific spatial patterns within constraints set by anatomical studies and the known characteristics of olfactory functioning.

What does the resulting model have to tell us about perceptual processing? Whenever an animal sniffs the process of transduction occurs: molecules of the odorant fall onto the cilia of olfactory receptor cells and excite a small subset of receptors in the nasal mucosa that are selectively sensitive to the odorant. These receptor cells send action potentials into the olfactory bulb that then excite the large projection neurons known as mitral cells. This is all that takes place in the olfactory system of unmotivated subjects. However, when a novel odorant is presented under reinforcement more takes place. In a series of sniffs synaptic strengths change in the bulb as the mitral cells, synaptically linked by bidirectional synapses, are co-activated pairwise by the CS + . This is a general form of the Hebb rule of synaptic learning (Hebb 1949). When many interconnected neurons within a mass of neurons fire together in pairs over repeated stimuli, the selectively co-activated neurons are joined together into a network of strengthened connections. This set is called a nerve cell assembly (NCA) and it is thought to be the basis for perception in the nervous system.

An essential feature of the model is the recognition that the formation of a NCA is not the whole story of perception, not even within the olfactory bulb (Skarda & Freeman 1988). With input, and under reinforcement, feedback interaction spreads excitation throughout the bulb. Other mitral cells that have been excited by background odorants, although not under reinforcement, excite the bulb further thereby increasing the sensitivity of bulbar neurons to each other. When a critical threshold is reached a further step occurs. This is a state change to a globally distributed pattern of activity in which all bulbar neurons take part, not just the limited number of neurons forming the NCA. The model postulates that the NCA operates at the moment of choice when the surge of input to the bulb strongly forces it away from its rest state to a new activity pattern. The NCA serves to amplify the small input received on any given inhalation, to produce a locally distributed low density activity pattern in interaction with the stimulus, and it provides the mechanism for mediating the emergence of the globally distributed, odor-specific activity pattern at the point of state change. This state change leads the entire bulb from apparently random activity to a globally coherent burst of oscillatory

activity that is passed on to the rest of the brain and that is hypothesized to constitute odor perception and recognition in the bulb.

This entire process is self-organized: with each inhalation after learning, and in the presence of the reinforced odorant, this more ordered global state emerges via an internally generated process from the less order background state, disappearing with each exhalation. A separate spatial pattern of oscillatory behavior forms for each odor given under reinforcement. When the reinforcement contingency is changed in respect to any one odor or when a new odorant is added to the animal's repertoire, all the globally distributed spatial patterns undergo change. These changes do not occur in the bulb if there is no reinforcement or if the newly learned stimulus is not olfactory but visual or auditory.

This picture of perceptual processing departs significantly from two important explanatory tools of modern physiological research. First, the data and the resulting model indicate that while the activity of single neurons appears to be largely unpredictable and noisy, networks and ultimately global assemblies of neurons cooperate to produce coherent, distributed patterns of activity that can be reliably correlated with particular stimuli. Perception in the olfactory system is expressed in the cooperative activity of masses of neurons and not in the activity of a favored few. This finding represents a break with one of the foundational concepts of contemporary research on the nervous system, the «neuron doctrine», the central assumption of which is that perception is caused by «the activity of a rather small number of neurons selected from a very large population of predominantly silent cells» (Barlow 1972). Olfactory data, along with recent data on neural dynamics in the visual cortex (Freeman & van Dijk 1987), contradict the single unit view of perceptual dynamics in the brain.

Second, the data contradict the classical model of physiological functioning developed to explain reflex behaviors (Sherrington 1906; Pavlov 1927), and the feedforward models experimentalists have used to explain how a conditioned stimulus will elicit a conditioned response. The basic assumption of reflex theories is that all behavior can be explained as the sum of responses to stimuli, however complex. According to this view, perception is a response, a reaction to receptor input, it is caused by the stimulus. But this is not an accurate physiological model of perceptual processing in the bulb. Causal impact of the stimulus on receptor cells is necessary for perceptual processing, but it does not constitute or even, as is often supposed, inaugurate perception. Perception begins *within* the organism with inter-

nally generated neural activity that, by re-afference, lays the ground for the processing of future receptor input. Evidence for this is based on the fact that presentation of odorants to the receptor surface in unmotivated subjects does not lead to any observable changes in the bulb; no NCA is formed and no global odor-specific activity patterns form. The brain itself internally initiates activity patterns that determine which, if any, receptor input will be processed by the brain. And after receptor activity is initiated, the brain itself produces self-organized patterns of activity that are not deterministic responses to stimuli; input de-stabilizes the olfactory bulb in such a way that internal patterned activity is released or allowed to develop (Freeman 1983). The story of perception cannot be told in terms of feed-forward causation because such a story overlooks the role of internally generated (self-organized) neural activity.

Perception, on this model, is interaction with the environment that is initiated by the organism. Olfactory data indicate that neural ensembles have evolved the capacity to select from the environment stimuli to which they will react, while neglecting a whole range of others. Perception is this process of selective interaction: «it is the organism itself... which chooses the stimuli in the physical world to which it will be sensitive» (Merleau-Ponty 1942). In this respect the neural system differs from all other physical systems, biological and nonbiological, with which we are familiar (Skarda & Freeman 1987). In other self-organizing physical systems we find the emergence of more ordered states from less ordered initial conditions, e.g., the formation of drops of water from a leaky faucet or the development of structure in a fertilized ovum, but only brains select the environment to which they will respond, thus providing the basis for what we term adaptive behavior.

What happens in the brain is about interaction. Olfactory research shows that there is no dependence of patterned activity in the bulb on odors presented to unmotivated subjects. Further, neither NCA formation nor globally distributed patterned activity take place in the absence of receptor activity. The networks and global ensembles that constitute perception in the olfactory system are interactive phenomena. Patterned activity correlates best with reliable forms of interaction in a context that is environmentally and behaviorally co-determined. This is the basic picture of perceptual processing that emerges from olfactory research.

3. *Connectionist Models of Perception*

Connectionist systems also use self-organized dynamics to model perceptual processing. Such systems can be very generally characterized as distributed systems made up of interacting simple elements. These systems do not use conventional computer programs that manipulate symbols according to rules. Instead, a perceptual state results by selectively strengthening and weakening the connection strengths between units that comprise the network in a highly parallel fashion. Once the network is activated by input to receptor level units, each unit computes its own level of activity in terms of input from all other units. The globally distributed patterned activity that results from these simultaneous, independent, parallel computations constitutes the state of the system at that moment.

In connectionist systems physical order (e.g., perceptual processing) results solely in virtue of the cooperative, internally generated activity taking place in assemblies of very simple elements. No single element in the system controls the others, none of the units represent or stand for anything, no program-specified rule tells the elements how to behave, and no symbols or internal representations are manipulated. These models are, from the biological perspective, more realistic than previous, symbol-based computational models and have succeeded at producing whole classes of behavior that eluded models using symbols and rules. Hence the interest they have recently generated is enormous, both in neuroscience and artificial intelligence circles.

But it is important for neuroscientists to note that there are respects in which the neural dynamics of perceptual processing differ from connectionist models. For one thing, connectionist systems still model perception as a reactive process, albeit a self-organized one; they encourage us to view network activity as a function of the causal impact of stimuli on receptor units. Connectionist systems use self-organized dynamics in the formation of the neural network, but the process of neural network formation is triggered by receptor input. The data from olfactory functioning indicate that this whole process can only take place in virtue of the prior existence of an internally generated activity state that is overlooked by connectionist systems and results in a mischaracterization of perceptual processing from the neurophysiological perspective. (This is not to say that connectionist research aims or should aim at producing neurophysiologically realistic

models of perception. The point is simply that neuroscientists and others interested in understanding how brains work ought to take such models with a grain of salt.)

Second, connectionist models indicate that the basic mechanism for perceptual processing is the formation of neural networks. As we have seen the model based on olfactory research argues that neural network activity is only part of the story of perception. Neural net formation and activity in this network constitute an intermediate stage of self-organization that is sandwiched between two other globally distributed self-organized processes: (1) an internally generated, stimulus-independent, prior activity state that sets the stage for processing of receptor input, and after receptor activation, (2) the stereotypical, odor-specific patterned activity that constitutes odor perception for the bulb. The formation of the neural net is a crucial step in perceptual processing, but it is only part of the story (Skarda & Freeman 1988).

There is a further important difference between the mechanisms used by the brain for perceptual processing and those used by connectionist systems. This difference is essential for evaluating the claim that perception consists in the construction and manipulation of internal representations. The distinction requires that we recognize that there are at least two classes of models within connectionist research (Freeman & Skarda 1988a).

The first class of models are parallel distributed processing models or PDP models (Hinton 1985; Rumelhart et al. 1986). Such systems rely on feed-forward connectivity and what is called back propagation for error correction. Most importantly, PDP systems have their «goals» externally imposed: they require a «teacher» or set of correct answers to be introduced by the system's operator, paradigmatic patterns (predetermined threshold values) against which the output of the system is corrected via back propagation and error correction. These mechanisms are important because although the teacher is not contained in a program, it nonetheless serves the same function that symbols do in conventional computers. For our purposes this means that perceptual processing in PDP models can be interpreted as a matching process in which a stimulus-evoked activity pattern is compared and adjusted with reference to an internally stored paradigm, model, or representation of the object perceived.

A second class of connectionist systems are self-organized dynamical (SD) models (Amari 1983; Anderson et al. 1977; Freeman

1987; Grossberg 1981; Hopfield 1982; Kohonen 1984). Both classes of connectionist systems use self-organizing dynamics and parallel, distributed processing, but SD systems replace the back propagation and error correction of PDP models with dense local feedback. As a result of this innovation, no teacher is required to predetermine values the network should achieve; instead, the network organizes itself from within to achieve the same ends. In such models the internal dynamics of the system are more truly described as self-organized and do not encourage us to look for an internal counterpart of the object perceived, or to view perception as a matching process. Instead, input destabilizes the system which responds by internally generating patterned activity via dense local feedback.

The dynamics of this second class of connectionist models, relying as they do on dense local feedback, resemble those used by brains, thus reinforcing the view that perception is mischaracterized when described as a process using internal representations, or as template matching. The nodes in connectionist networks do not represent anything and in SD connectionist systems no teacher functions like a representation either. If the brain resembles SD systems, and the evidence indicates that it does (Skarda & Freeman 1987), then it seems unconvincing to accept the representationalist interpretation.

But there are further grounds to reject representationalism. The position as it is usually stated holds that the goal of perceptual processing is the creation of an internal model of an object and its features. Given this we might expect that simple presentation of that object would lead to the formation of stimulus-specific activity patterns in the brain. But this does not happen when an animal perceives. As previously discussed, presentation of odorants to an unmotivated animal has no effect, i.e., no NCA forms and no odor-specific activity patterns are generated. It seems reasonable to conclude that not only is perception not a reaction to an object, it isn't just about objects either; rather, perception is about objects in interaction with brains that first internally generate an activity state that permits receptor activity to «take hold». This conclusion is supported by the further finding that the patterns that do form in motivated subjects and that we can correlate with specific odors are dependent on behavioral response as well as context. If we change the reinforcement contingency or introduce new odors to the repertoire all of the previously recorded activity patterns undergo change. If these patterns were representations of the stimulus and its features we would not expect

such changes to occur on presentation of one and the same odorant. But do these findings simply force us to elaborate our definition of «object», so that by it we mean past experiences and context as well as the object perceived? I do not think so.

What happens in the brain can be explained without recourse to representations. Motivation leads to the creation of a self-organized internal state that destabilizes the system in such a way that it becomes ready to respond to a particular class of stimulus input in a sensory modality (Skarda & Freeman 1987). This class of stimuli may or may not have been experienced previously and may be quite general, but once it is received it, in turn, sets up conditions such that the system will generate a new form of interactive behavior to cope with the constraints imposed by new circumstances and previous experiences (Skarda 1986). Perception is an interactive process of destabilization and re-stabilization via self-organized dynamics. Each state change requires a qualitative (parametric) change in the system, and not merely a change in its input.

There is nothing intrinsically representational about this process, at least not until the observer enters the picture. The experimenter's viewpoint is the factor that requires that conclusions be drawn about what the observed activity patterns represent to the subject. Perhaps representation-based accounts of perception have simply incorporated this external viewpoint in their model of perceptual processing. This is both unnecessary, as evidenced by self-organized dynamical connectionist systems and brain dynamics, and misleading in our attempts to understand perception (Freeman and Skarda 1988b). It is misleading because instead of looking at neural dynamics in its own right, we focus on, e.g., pattern invariance or «feature detectors». Olfactory research gives us no evidence that internal representations exist. But even if internal representatives of the environment were found in the brain, we would be no closer to understanding *brain function*. We would still not know how the patterns were generated from less ordered initial conditions, their effects on neurons to which they transmit, the principles of neural functioning they embody or the kinds of structural changes involved. In the end, representationalism gives researchers the illusion of having discovered something about the brain because something is known about the objects which it claims are represented inside. This illusion impedes, it does not foster, advances toward understanding brain function.

4 Conclusions

Until recently, it was fashionable to claim that the study of cognition had nothing to learn from neuroscience. From our vantage point today, with the advent of connectionism and renewed interest in neural networks, such a view seems misguided and naive. Proponents of the view were miserably uninformed about neuroscience: they misrepresented the scope and character of data, and had little idea therefore of the range of questions upon which the data could be brought to bear (Churchland 1986; Skarda 1986).

But it is one thing to argue that neuroscience is relevant, and another to show how its findings can be applied. Many cognitive theorists today are convinced that neuroscience has a valuable contribution to make, but in practice most carry out their work with little more than a passing reference to the relevant neuroscientific literature. The general consensus seems to be that the neuroscience that will prove relevant to cognitive science is not the one that presently exists, but a future neuroscience that has progressed far beyond anything we now know.

This paper had one objective: to explain how the data on neural processing in the olfactory system force a reformulation of our understanding of perception. In the process I hope to have persuaded the reader that neuroscience is relevant to cognitive science, and that it is relevant now. Data from olfactory processing forced us to reconsider previously held views of perceptual processing: perception is not reaction to a stimulus, it does not require appeal to a nonbiological element or force, nor is it internal representation. Perception is not the kind of process we imagined it to be. Connectionist models, for all the interest they engender, have not forced the kind of radical reappraisal of perceptual processing that olfactory research forces upon us. Perhaps this was to be expected: formal models, after all, reflect our current understanding and present theories of physical functioning. While connectionists were constructing their models, brains were busy doing the unexpected. Perhaps, this is the real lesson of neuroscience today for cognitive science, and the secret behind its unique contribution.

BIBLIOGRAPHY

- AMARI, S. (1983) Field theory of self-organizing neural nets. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13, No. 5: 741-48.
- ANDERSON, J., SILVERSTEIN, J., RITZ, S. & JONES, R. (1977) Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review* 84: 413-51.
- BARLOW, H. (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1: 371-94.
- CHURCHLAND, P. (1986) *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. MIT Press/Bradford.
- FREEMAN, W. (1983) Dynamics of self-organizing neural systems in attention and perception of rabbits: A reflection of Katchalsky's contribution. In *Biological Structures and Coupled Flows*, eds. A. Oplatka & M. Balaban. Academic Press.
- 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. & SKARDA, C. (1985) Spatial EEG patterns, nonlinear dynamics and perception: The neo-Sherringtonian view. *Brain Research Reviews* 10: 147-75.
- FREEMAN, W. & SKARDA, C. (1988a) Mind/brain science: Neuroscience on philosophy of mind. In *John Searle and His Critics*, eds. E. LePore & R. van Gulick. Blackwell.
- FREEMAN, W. & SKARDA, C. (1988b) Representations: Who needs them? In *Proceedings 3rd Conference on the Neurobiology of Learning and Memory* (forthcoming).
- FREEMAN, W. & VAN DIJK, B. (1987) Spatial patterns of visual cortical fast EEG during conditioned reflex in a rhesus monkey. *Brain Research* 422: 267-276.
- 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.
- HINTON, G. (1985) Learning in parallel networks. *Byte* 10:265.
- HOPFIELD, J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* 79: 2554.
- KOHONEN, T. (1984) *Self-Organization and Associative Memory*. Springer Verlag.
- MERLEAU-PONTY, M. (1942) *La Structure du Comportement*. Presses Universitaires de France.
- PAVLOV, I. (1927) *Conditioned Reflexes*. Oxford University Press.
- RUMELHART, D., McCLELLAND, J. & PDP RESEARCH GROUP (1986) *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*. 2 Vols. MIT Press/Bradford.
- SKARDA, C. (1986) Explaining behavior: Bringing the brain back in. *Inquiry* 29: 187-202.
- SKARDA, C. & FREEMAN, W. (1987) How brains make chaos in order to make sense of the world. *Behavioral and Brain Sciences* 10: 161-173.
- SKARDA, C. & FREEMAN, W. (1988) Research of neural dynamics: Implications

for models of learning and memory. *Systems with Learning and Memory*. Springer Verlag.

SHERRINGTON, C. (1906) *The Integrative Action of the Nervous System*. Yale University Press.