J. Theoret. Neurobiol. **1**, 286–369 (1982)
Printed in Australia

A PSYCHOPHYSIOLOGICAL THEORY OF REINFORCEMENT, DRIVE, MOTIVATION AND ATTENTION

STEPHEN GROSSBERG*
Center for Adaptive Systems,
Department of Mathematics, Boston University,
Boston, Massachusetts 02215, U.S.A.

(Received February 16, 1982)

Abstract

This article derives a real-time theory of motivated behavior and presents some of its physiological and pharmacological correlates. The theory mechanistically explicates instrumental concepts such as reinforcement, drive, incentive motivation, and habit, and describes their relationship to cognitive concepts such as expectancy, competition, and resonance. The theory shows how a real-time analysis of an animal's adaptive behavior in prescribed environments can disclose network principles and mechanisms which imply a restructuring and unification of the data in terms of design principles and mechanisms rather than the vicissitudes of experimental methodology or historical accident. A comparative analysis and unification of other theories is then possible, such as the classical theories of Hull, Spence, Neal Miller, Estes, Logan, Livingston, and John. The data which are discussed include overshadowing and unblocking; suppression by punishment; reinforcement contrast effects; hypothalamic self-stimulation; differential effects of drive, reinforcement, incentive motivation, expectancies, and short-term memory competition on learning rate, behavioral choice, and performance speed; the role of polyvalent cortical cells, multiple sensory representations, recurrent on-center off-surround neocortical and paleocortical interactions, hippocampal–hypothalamic, medial forebrain bundle, and thalamocortical interactions on motivated behavior; effects of drugs like chlorpromazine, reserpine, monoamine oxidase inhibitors and amphetamine on instrumental behavior. Of special interest are network 'hippocampal' computations that are suggested to accomplish several distinct roles: influence transfer of short-term to long-term memory both directly and indirectly, directly by triggering conditioning of conditioned reinforcers, indirectly by generating positive attentional feedback to neocortical polyvalent cells; and influence the organization of a motor map which controls approach and avoidance behavior by eliciting motivationally biased signals to this motor mapping system.

* Supported in part by the National Science Foundation (NSF IST-80-00257) and the Air Force Office of Scientific Research (AFOSR 82-0148).

1. Introduction

The phenomena of classical and instrumental conditioning are wonderfully varied, interesting, and confusing (Bolles, 1967; Brush, 1971; Campbell and Church, 1969; Estes, 1969; Gilbert and Sutherland, 1969; Honig, 1966; Mackintosh, 1974; Pearce and Hall, 1980; Rescorla and Wagner, 1972; Tapp, 1969; Wagner, 1978). It has often seemed nigh impossible to state a law, or generalization, in this area that does not admit important exceptions. For example, are classical conditioning and instrumental conditioning mechanistically independent? If not, exactly how are they related? Do specific drives really exist, such as a hunger drive? Do generalized drives exist, such as an exploratory drive? If so, why do both types of drives exist? Are drives independent, such as hunger and thirst, or hunger and exploratory drive? What is the difference between a drive and an incentive? What is the difference between a drive and a habit? Do drives have some properties of habits? In particular, can drives act as stimuli that can be associatively joined to responses, or do they act simply as sources of energy? Is reinforcement due to drive reduction? If so, what drive is reduced to reinforce language behavior? If not, why does reinforcement often seem to be related to drive reduction? The list of questions goes on and on. After a while, one can become exhausted by the sheer variety of concepts, and by the seemingly endless introduction of exceptional rules or subtle distinctions to handle difficult special cases.

Because these difficulties are so formidable on the level of consensual language, we need to find a formal language which is powerful enough to describe instrumental concepts and mechanisms without ambiguity. A method must be found whereby this formal language can be constructed. Such a method has been gradually developed over the past two decades. The method starts by identifying environmental pressures to which an organism must adapt in order to survive. In other words, the theory is derived from real-time constraints on the self-organization (learning, development) of individual behavior. The method shows how each environmental problem is solved by a principle of behavioral organization. These principles are realized by mathematical laws which embody the principles in the most parsimonious, or minimal, way. Once a minimal solution is found, one can more readily classify individual and species differences as variations on the minimal theme. In every case studied to the present, these mathematical realizations can be interpreted as neural networks.

The derivation of these principles and laws can be achieved using thought experiments which show us, in simple stages, how prescribed environmental pressures force adaptive designs on the behaving brain. I contend that various other psychological theories have failed to derive significant conclusions about brain design because they have ignored real-time constraints on the self-organization of individual behavior. The concepts that arise from this procedure are, not surprisingly, related to traditional ideas, but they diverge from traditional ideas in crucial ways that allow us to penetrate into areas where the traditional notions are misleading or too vague to follow.

This procedure eventually leads to a principled reorganization of the data. For example, our first thought experiment is about a dilemma concerning classical conditioning (Grossberg, 1971*a*). Namely, how does a network learn associatively despite the fact that the time intervals between successive CS and UCS presentations can vary across learning trials? This seemingly innocent question forces us into explicit mechanisms of instrumental conditioning, into a role for cognitive processing in the direct evaluation of reinforcement, and to the threshold of studies about short term memory, attention, and the development of cognitive codes. Herein I follow the path by which these mechanisms were historically derived, since it provides a convenient route through the data along a gradually rising conceptual pathway. Each constraint along this pathway provides us with a necessary constraint on network design, which we then translate into a minimal realization before classifying related possibilities. Similar mechanistic conclusions can be derived from a thought experiment about cognitive development (Grossberg, 1980*a*, 1982*a*). This multiplicity of derivations, always leading to conclusions that support and sharpen previous results, endows the theory with an aura of conceptual coherence and robustness.

The historical procedure is not without its expository difficulties. Each stage of the derivation sheds new light on a variety of nontrivial data. However, in a living organism, often mechanisms other than the ones that have just been derived are also at work. My choice is whether to defer all interpretations until the whole theory is derived, or to provide data markers along the way using experiments that emphasize the mechanisms then being discussed. I have chosen the latter procedure and along the way will try to direct the reader to the related mechanisms too.

For example, I was led to mechanisms for reinforcement and incentive motivation before being driven by the theory to consider expectation mechanisms. In the theory, these two classes of mechanisms often closely interact, but are distinct. In the data, it is often hard to tease them apart. Thus Mackintosh (1974, p. 233), after a sophisticated data analysis, wrote that

> Stimuli associated with reinforcements do not motivate instrumental responses; they may become established as goals for instrumental responses. Their effect on instrumental behavior therefore may be similar to that of unconditional reinforcers; they may, in other words, serve as conditional reinforcers.

I will also argue that stimuli which are associated with reinforcements can serve as conditioned reinforcers. However, I will argue, contrary to Mackintosh, that such stimuli can motivate instrumental responses and that, although they may become established as goals, this goal property is, strictly speaking, related to expectancy mechanisms and not to reinforcer properties *per se*.

To understand the distinction between Mackintosh's claim that reinforcers do not motivate instrumental responses, and my claim that they do, one must understand what position Mackintosh is arguing against; namely, the classical idea that motivation acts in a nonspecific fashion. The theoretical development in this paper argues against the classical position, but also leads to a mechanism of motivation that is important in reinforced behavior. The theory also shows how the reinforcing

properties of stimuli can easily be confused with their motivating properties. It hereby suggests how Mackintosh could be led by considerations of parsimony to eliminate motivation from his discussion. I will argue that the two concepts are really fundamentally distinct, and will attach them to distinct psychophysiological parameters.

Once we embrace the evolutionary method, we must be prepared to organize our mechanistic understanding into a succession of conceptual stages. The article is structured as such a succession. This procedure is inherent in the evolutionary method, and is the price we pay for understanding very well, sometimes painfully well, just what each stage's organizational principles do and do not imply.

2. Buffer the learning cells

Before beginning the thought experiment, let us put it into perspective with some introductory remarks It is obvious that cells which are capable of learning should be buffered against being activated except by appropriate inputs. In particular, during adult human behavior, cells near the sensory and motor periphery should not be capable of substantial learning. If adult retinas could learn, we would see a superposition of all the visual scenes that occurred for hours past. If adult motor neurons could learn, our next motion would be an inaccurate weighted average of all our recent motor commands. To prevent this, the cells that are most capable of learning will, in the adult, be found away from the sensory or motor periphery, where they are carefully surrounded by protective networks. This is one reason why learning cells have either been hard to isolate by blind electrode penetrations, or where cells have been reliably isolated, they often have disappointing learning capabilities (Hoyle, 1977; Kandel, 1976; Morrell, 1961).

3. A digression on classical conditioning

The thought experiment demonstrates that classical and instrumental conditioning share certain mechanisms in common. These mechanisms embed, or buffer, the cells capable of learning in a network that prevents the cells' activation except under appropriate circumstances. The thought experiment builds upon prior work which derives laws of associative learning from the simplest concepts of classical conditioning; see Grossberg (1974, 1982*c*) for a review. In this work, laws for neural networks are derived from a real-time analysis of how pairing a conditioned stimulus (CS) with an unconditioned stimulus (UCS) on learning trials enables the CS to elicit a conditioned response (CR), or UCR-like event, on performance trials. The network dynamics are described by interactions between the *short-term memory* (STM) *traces* $x_i(t)$ of cell body populations v_i, and the *long-term memory* (LTM) *traces* $z_{jk}(t)$ of the axonal pathways e_{jk} from v_j to v_k, as in Figure 1. For present purposes, the exact form of these interactive STM-LTM laws is not important. What is important are two properties of these laws.

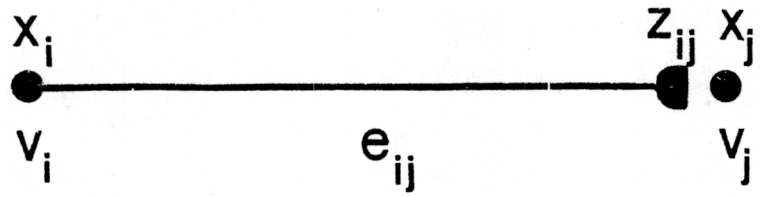

Fig. 1. Short-term memory traces (or potentials) x_i at cell populations v_i emit signals along the directed pathways (or axons) e_{ij} which are gated by long-term memory traces z_{ij} before they can perturb their target cells v_j.

A. The unit of LTM is a *spatial pattern.*

B. There exists a *stimulus sampling operation.*

By (A) I mean the following. Consider the network in Figure 2*a*. This network depicts the minimal anatomy that is capable of learning by classical conditioning. The population v_0 receives a CS-activated input. Population v_0 can thereupon emit signals along its axons e_{0i} whose terminals, or synaptic knobs, S_{0i} $(i = 1, 2, \ldots, n)$ abut on the UCS-activated populations $v_1, v_2, \ldots, v_n$. The LTM traces z_{0i} are computed at the synaptic knob terminals. Each z_{0i} computes a time average of the signal along e_{0i} multiplied by the STM trace x_i of v_i. In particular, z_{0i} cannot discern the value of x_i unless the axonal signal is positive. *Stimulus sampling* means that z_{0i} can only detect the effect of the UCS pattern on x_i at times when signals from v_0 reach the synaptic knobs S_{0i}.

The network of Figure 2*a* is called an *outstar* because it can be symmetrically redrawn as in Figure 2*b*. Property (A) means that an outstar can learn an arbitrary spatial pattern. A *spatial pattern* is a UCS to the cells $v_1, v_2, \ldots, v_n$ whose intensities have a fixed relative size through time; that is, the input $I_i(t)$ to v_i satisfies $I_i(t) = \theta_i I(t)$, $i = 1, 2, \ldots, n$. The constants, or 'reflectances', θ_i are nonnegative and are normalized such that $\sum_{k=1}^{n} \theta_k = 1$ to achieve the convention that $I(t)$ is the total UCS input; viz. $I(t) = \sum_{k=1}^{n} I_k(t)$. The outstar can learn the pattern weights $\theta = (\theta_1, \theta_2, \ldots, \theta_n)$ at a rate that depends upon the size of the CS input $I_0(t)$ and the total UCS input $I(t)$ (Grossberg, 1970*a*).

The *stimulus sampling probabilities* of an outstar are the relative LTM traces

$$Z_{0i} = z_{0i}\left(\sum_{k=1}^{n} z_{0k}\right)^{-1}. \tag{1}$$

As CS–UCS pairing takes place, the functions Z_{0i} approach the values θ_i, respectively. During later performance trials, a CS input to v_0 creates equal signals in the e_{0i} axons. These signals are gated, or multiplied, by the LTM traces z_{0i}. Since each z_{oi} is proportional to θ_i, the gated signal to v_i is also proportional to θ_i. The

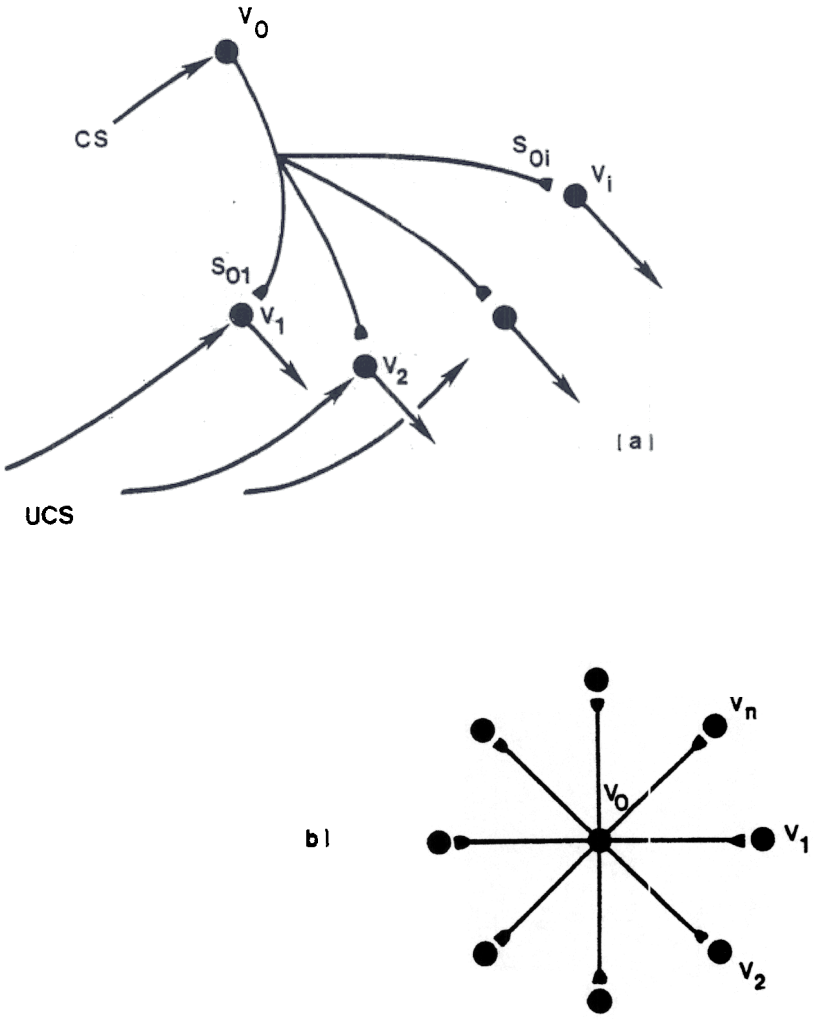

Fig. 2. In (*a*) the conditioned stimulus (CS) activates population v_0 which thereupon sends sampling signals to the unconditioned stimulus (UCS) activated populations $v_1, v_2, \ldots, v_n$; in (*b*) the *outstar* is the minimal network capable of classical conditioning.

CS hereby elicits responses in the STM trace x_i that are proportional to θ_i. In short, after CS–UCS pairing, the CS can reproduce the pattern θ.

Stimulus sampling can be described as follows. The stimulus sampling probabilities Z_{0i} can change *only* when signals from v_0 reach the synaptic knobs S_{0i}. Unless the CS perturbs these knobs, their LTM traces cannot 'see' what UCS patterns are received at the cells $v_1, v_2, \ldots, v_n$.

These simple ideas about classical conditioning can be generalized to prove a universal theorem about associative learning. The universal theorem guarantees

unbiased spatial pattern learning by arbitrarily many, simultaneously active sampling populations that are activated by arbitrary continuous data preprocessing in an essentially arbitrary anatomy (Grossberg, 1969*a*, 1971*b*, 1972*a*). The same laws also, for example, imply many properties and predictions of serial learning, paired associate, and free recall data (Grossberg, 1969*b*, 1978*a*; Grossberg and Pepe, 1971). In the present article, I will develop implications of these laws that are based on properties (A) and (B) above.

4. Motor synergies as evolutionary invariants

Although I will not dwell on these applications and generalizations of the associative learning laws, the reader should realize that the outstar is a general purpose pattern learning device. To illustrate this fact, suppose that the pattern weights θ describe the relative activities of motor control cells, and that these constant relative activities are transmuted into fixed relative rates of muscle contraction across their controlled muscle groups. Then the fact that an outstar can learn and later perform a spatial pattern without destroying the memory of the pattern means that a motor synergy can be learned and stably performed by a single command population v_0. The fact that the same pattern θ can be performed by different CS sampling signals means that performance of the motor synergy is effected by synchronous contraction of all the muscles. Distinct synchronous performance signals can alter the absolute muscle contraction rates through time, but preserve the relative contraction rates. Since the muscle positions before outstar read-out do not equal the terminal motor pattern encoded by θ, these invariant ratios will be easier to measure nearer to the end than the beginning of the synergetic motion. Analogous properties have recently been reported during motor performance (Kelso *et al.*, 1979; Soechting and Lacquaniti, 1981). From the perspective of the present theory, these properties reflect an invariant of the learning or evolutionary process; namely, that outstars encode pattern ratios which are left invariant by synchronous performance signals.

5. A thought experiment: the synchronization problem of classical conditioning

Our thought experiment will be based on obvious real-time constraints which classical conditioning imposes upon a behaving individual. The fact that the constraints are obvious does not mean that they are trivial. In the present case, it means they are so ubiquitous and we have adapted so well to them that they seem obvious. The very obviousness of these constraints gives force to our argument.

The two main constraints are the following (Grossberg, 1971*a*).

C. The time intervals between CS and UCS presentation on successive learning trials can differ; and

D. The CS alone can elicit a CR on performance trials.

Postulate (C) describes the obvious fact that successive stimulus presentations under natural conditions in real-time are not always perfectly synchronized; postulate (D) describes the outcome of classical conditioning, and simply asserts that this simplest example of associative learning is possible. We now show that to satisfy these postulates in a world wherein events continually buffet our senses, and wherein our long-term memories are spatially, albeit nonlocally, coded requires additional network structure.

To see this, we observe a continual stream of patterns from the viewpoint of an outstar. We ask, how can an outstar learn anything at all if unsynchronized patterns continually flow by? In particular, suppose that the outstar $\mathcal{O}_1$ attempts to learn a prescribed pattern $\theta^{(1)}$ in a sequence $\theta^{(1)}, \theta^{(2)}, \theta^{(3)}, \ldots$ of spatial patterns by practising as the sequence is presented on successive learning trials. Denote $\mathcal{O}_1$'s sampling population, or source, by $v_1^{(2)}$ and $\mathcal{O}_1$'s sampled populations, or border, by the field $\mathcal{F}^{(1)} = \{v_1^{(1)}, v_2^{(1)}, \ldots v_n^{(1)}\}$. If postulate (C) holds, then the time lag between the CS that excites $v_1^{(2)}$ and the onset of the UCS sequence $\theta^{(1)}, \theta^{(2)}, \theta^{(3)}, \ldots$ that perturbs $\mathcal{F}^{(1)}$ can be different on successive learning trials. If $v_1^{(2)}$ fires whenever the CS occurs, then $\mathcal{O}_1$ can sample a different pattern $\theta^{(k)}$ on every learning trial. $\mathcal{O}_1$ will consequently learn an average pattern that is derived from all the sampled patterns; i.e., 'noise'. To avoid this catastrophe, $\mathcal{O}_1$ must know when to sample the 'important' pattern $\theta^{(1)}$. Somehow the onset of sampling by $v_1^{(2)}$ and the arrival of the UCS at the field $\mathcal{F}^{(1)}$ of sampled cells must be synchronized so that $\mathcal{O}_1$ can sample $\theta^{(1)}$, and only $\theta^{(1)}$, on successive trials.

How can the onset of $v_1^{(2)}$ sampling be synchronized to occur a fixed time before the UCS arrives at $\mathcal{F}^{(1)}$ if the CS and UCS onset times are themselves unsynchronized? This can only happen, in principle, if several properties are imposed.

First, the CS itself must be insufficient to elicit a sampling signal from $v_1^{(2)}$. Second, the UCS must let $v_1^{(2)}$ know when it will arrive at $\mathcal{F}^{(1)}$ by sending a signal to $v_1^{(2)}$. Third, $v_1^{(2)}$ must be prevented from eliciting a sampling signal unless large CS *and* UCS signals converge simultaneously at $v_1^{(2)}$. In other words, $v_1^{(2)}$ should not fire at all unless it represents the CS and should not fire until the correct time before the UCS arrives at $\mathcal{F}^{(1)}$. In particular, if the CS input arrives so long before the UCS that its signal to $v_1^{(2)}$ decays before the UCS signal reaches $v_1^{(2)}$, then $v_1^{(2)}$ cannot fire. Fourth, the UCS signal must arrive at $v_1^{(2)}$ before the UCS pattern activates $\mathcal{F}^{(1)}$, since $v_1^{(2)}$ must be able to send a signal to $\mathcal{F}^{(1)}$ in time to sample $\theta^{(1)}$. In other words, the UCS activates a bifurcating pathway. One branch in the pathway *arouses* $v_1^{(2)}$; that is, it gets $v_1^{(2)}$ ready to sample the UCS that will soon perturb $\mathcal{F}^{(1)}$. The other branch delivers the UCS pattern to $\mathcal{F}^{(1)}$ a little while later (Figure 3). Fifth, the UCS does not know to which CS it will be paired in a given experiment. It could be paired with any CS! The above argument holds for every CS with which the UCS can possibly be associated. Thus the UCS must be able to arouse all of the sampling cells that these CS's activate; namely, the whole field $\mathcal{F}^{(2)} = \{v_1^{(2)}, v_2^{(2)}, \ldots, v_N^{(2)}\}$ of CS-activated sampling cells. Thus the

UCS *nonspecifically* arouses the entire field $\mathscr{F}^{(2)}$ just before it delivers its pattern to $\mathscr{F}^{(1)}$ (Figure 4).

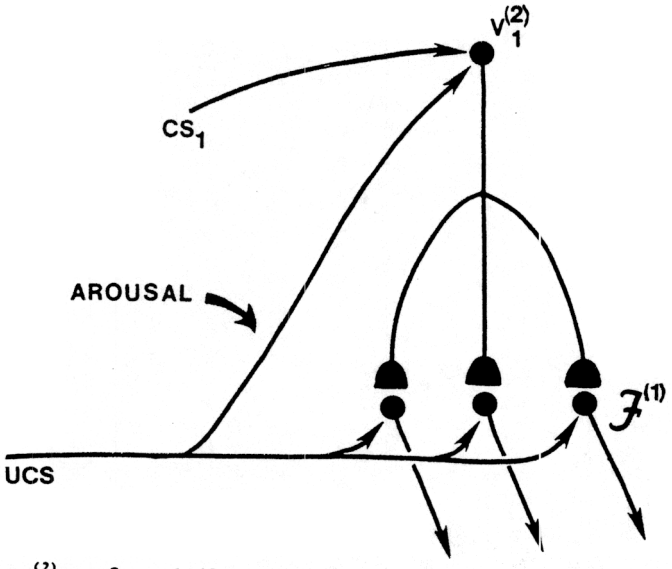

Fig. 3. Population $v_1^{(2)}$ can fire only if CS and UCS-arousal signal simultaneously converge upon it. The UCS input bifurcates to deliver a UCS pattern at $\mathscr{F}^{(1)}$ and arousal to sampling cells like $v_1^{(2)}$.

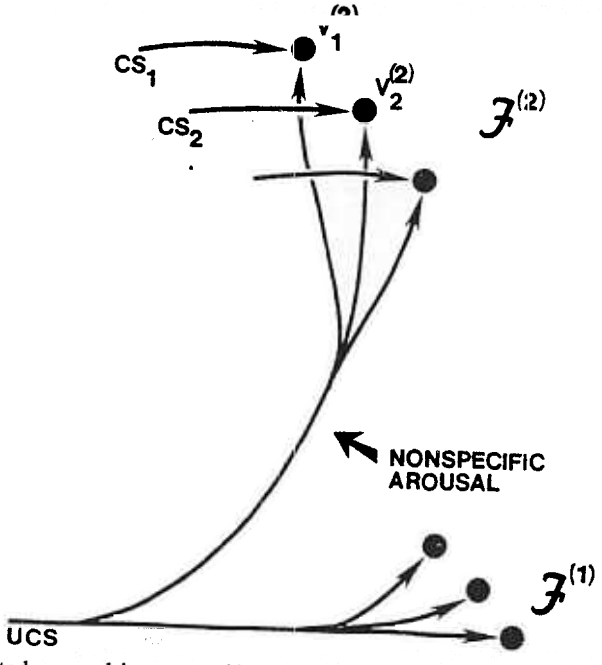

Fig. 4. UCS-activated arousal is nonspecifically delivered to all the sampling cells $\mathscr{F}^{(2)}$ because it cannot be known *a priori* which CS–UCS association will be imposed by the environment.

In summary, simultaneous convergence of the CS input and the UCS nonspecific arousal at a sampling cell are needed to fire this cell. This mechanism synchronizes the onset of CS-activated sampling signals from $\mathscr{F}^{(2)}$ and the arrival of UCS patterns at $\mathscr{F}^{(1)}$ on successive learning trials. Synchronization is a necessary condition in order for practice on successive trials to avoid massive associative confusions among spatially encoded patterns that stream into the network through time.

6. Some experimental connections

Before continuing with the thought experiment and the consequent derivation of increasingly precise neural structure, let us realize that some basic psychophysiological facts are already coming into view.

(i) *Nonspecific arousal*

The thought experiment teaches us that sampling cells cannot fire unless they are nonspecifically aroused, even if they receive specific CS inputs. At least since the work of Moruzzi and Magoun (1949), it has been known that inactivity of nonspecific subcortical projection systems to the cerebral cortex can prevent the cortex from supporting conscious behavior. The field $\mathscr{F}^{(2)}$ of sampling cells become a rudimentary analog of the cortex in our thought experiment.

(ii) *Events have cue and arousal functions*

To organize the flood of data that followed Moruzzi and Magoun's study, Hebb (1955) suggested that every sensory event has two quite different effects: its *cue* function and its *arousal* or *vigilance* function. The cue function represents the information in the event that selectively guides behavior. The arousal function energizes the behavior. Hebb suggested that learning without arousal is not possible. Hull (1943) had earlier dichotomized information and energetic variables by distinguishing habit strength (${}_SH_R$) from drive (D). Hull suggested that drives energize habits via the multiplicative law ${}_SE_R = {}_SH_R \times D$ for reaction potential ${}_SE_R$. Actually, the distinction between information, or reason, and energy, or passion, is a very old one that was already embraced by the rationalists (Bolles, 1967) in their efforts to construct a comprehensive philosophical framework by which to understand human behavior. The distinction has even been a force guiding social policy as in Vienna during the time of Wittgenstein (Janik and Toulmin, 1973), where men were supposed to embody the principle of reason, and women the principle of passion that was considered to be destructive of reason. This belief was used to justify various unpleasant social policies. By contrast with the Viennese notion, the thought experiment requires both principles to compute the simplest memories, reasonable or not.

In Figure 4, the UCS has both a cue and an arousal function due to its bifurcating pathway, but the CS has only a cue function. Does this distinction say something basic about the difference between CS and UCS? The next section will suggest an answer.

(iii) *Polyvalent cells*

The thought experiment implies that the sampling cells can be polyvalent cells, or cells that are influenced by more than one modality. For example, if the CS is a tone and the UCS is a visual cue of food, then both auditory and visual cues will be needed to fire the corresponding sampling cells. In other words, the sampling cells only fire in response to the sum of CS and UCS inputs, and their firing patterns influence the occurrence or nonoccurrence of network learning. John (1966, 1967) has discovered neocortical cells that satisfy all of these properties.

(iv) *D.C. potential shifts*

Such workers as Rusinov (1953) and Morrell (1961) have shown that electrodes that induce anodal d.c. potential shifts of cortical tissue can augment conditioning within the cortex, whereas cathodal d.c. shifts tend to inhibit conditioning. These results suggest that the anodal d.c. shifts have effects on cortical activity that are analogous to nonspecific arousal.

7. Conditioned arousal

To continue the thought experiment, we now consider postulate (D). This postulate is that classical conditioning is possible, or that the CS alone can elicit a CR after conditioning occurs. In particular, after conditioning occurs, a CS input to $\mathscr{F}^{(2)}$ can elicit a sampling signal that reads out the learned pattern across $\mathscr{F}^{(1)}$. Postulate (D) forces us to consider the following paradox. If a cell in $\mathscr{F}^{(2)}$ can only be fired when a specific CS signal and a nonspecific UCS signal converge on the cell, then how does the CS alone fire the cell after learning occurs?

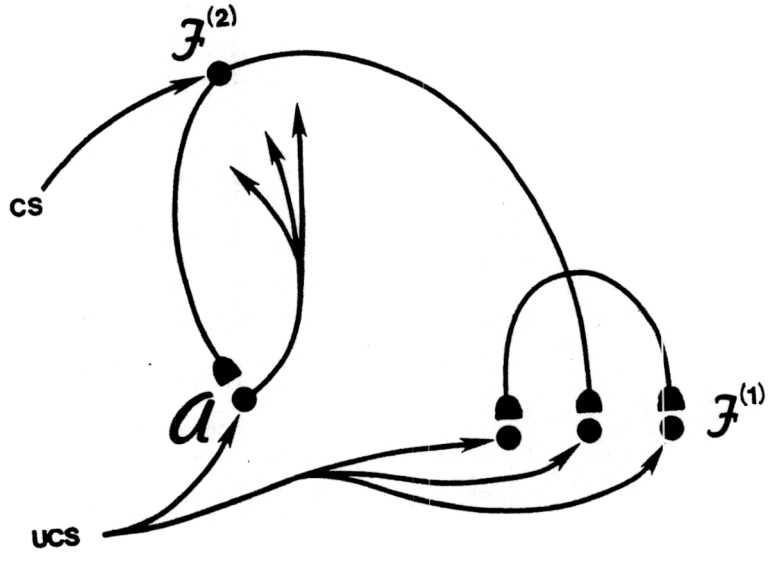

Fig. 5. On learning trials, the CS samples both the UCS pattern at $\mathscr{F}^{(1)}$ and the nonspecific arousal source $\mathscr{A}$.

In order for a cell $v_i^{(2)}$ to fire, it must receive simultaneous specific and nonspecific signals. However, only the CS is present on recall trials to activate both of these signals, since during recall trials, no UCS is presented to activate the nonspecific pathway. Thus the CS itself must activate both pathways. Before learning occurs, the CS does not have this capability. Somehow, as a result of pairing the CS and the UCS on learning trials, the CS gains control over the nonspecific pathway that is activated by the UCS. In other words, the nonspecific arousal pathway can be conditioned, and the CS can sample this conditionable pathway (Figure 5).

In summary, two conditioning processes occur in parallel during classical conditioning: (1) the CS, via $\mathscr{F}^{(2)}$, samples the cells in $\mathscr{F}^{(1)}$ that control the CR; and (2) the CS also samples the cells that control nonspecific arousal. Once the CS accomplishes both tasks, it can, by itself, fire cells in $\mathscr{F}^{(2)}$ that read out the learned CR pattern from $\mathscr{F}^{(1)}$.

An important implication of this argument is that there must exist cells, other than sampling cells $\mathscr{F}^{(2)}$ and sampled cells $\mathscr{F}^{(1)}$, that participate in classical conditioning. These are the arousal cells, denoted by $\mathscr{A}$ in Figure 5, at which the CS and UCS signals gain control of nonspecific arousal signals to the sampling cells.

8. Secondary reinforcers

The network in Figure 5 begins to explain how secondary reinforcers operate. The UCS in a given learning experiment might have been only the CS in a previous learning experiment. How is this possible? For example, consider an animal $\mathscr{S}$ at two successive stages, E_1 and E_2, of its development. At stage E_1, $\mathscr{S}$ salivates in response to the smell of food but not to visual presentation of food. After classical conditioning with CS = visual presentation of food and UCS = smell of food, the animal salivates when it sees the food. This ability characterizes stage E_2.

Now a second conditioning experiment is performed in which CS = ringing bell and UCS = visual presentation of food. Ultimately, $\mathscr{S}$ salivates when it hears the ringing bell. How does the visual presentation of food in the first experiment enable this event to become a UCS in the second experiment?

The network in Figure 5 suggests an answer. The UCS of the first experiment is a UCS because it controls both the nonspecific arousal pathway and a specific pathway. The CS of this experiment becomes a UCS by gaining control over the nonspecific arousal pathway as well as its specific pathway. We can now begin to see that activation of the nonspecific arousal pathway is closely related to the motivational properties of the UCS, and that any cue that can activate such a pathway acquires motivational properties. The synchronization property of classical conditioning has hereby begun to force us into basic mechanisms of instrumental conditioning.

9. Minimal network realization of conditioned nonspecific arousal

We now have enough information available to construct the minimal network capable of conditioning nonspecific arousal and using it to elicit overt behavior. We know that the CS, via cells in $\mathscr{F}^{(2)}$, can learn to activate nonspecific arousal

from $\mathscr{A}$. The arousal, in turn, is needed to elicit sampling signals to $\mathscr{F}^{(1)}$. We now show that the CS-activated cells that sample $\mathscr{S}$ and that receive feedback from $\mathscr{S}$ cannot be the same cells.

Figure 6 depicts the four general ways in which the CS can act. Figure 6*a* is impossible for the following reason. The CS-activated cells $\mathscr{F}^{(2)}$ cannot fire on recall trials unless they are aroused by $\mathscr{A}$. Cells in $\mathscr{A}$ cannot fire, however, unless they are activated by $\mathscr{F}^{(2)}$! Hence the cells that sample $\mathscr{A}$ cannot be the same as the cells that are aroused by $\mathscr{A}$.

Figure 6*b* tries to remedy this difficulty by expanding the cells $v_i^{(2)}$ of $\mathscr{F}^{(2)}$ into two successive stages $v_{i1}^{(2)}$ and $v_{i2}^{(2)}$ of processing, in which the first stage $v_{i1}^{(2)}$ excites the second stage $v_{i2}^{(2)}$, and in which one stage samples $\mathscr{A}$ and the other stage is aroused by $\mathscr{A}$. In Figure 6*b*, however, the connections to $\mathscr{A}$ are in the wrong order.

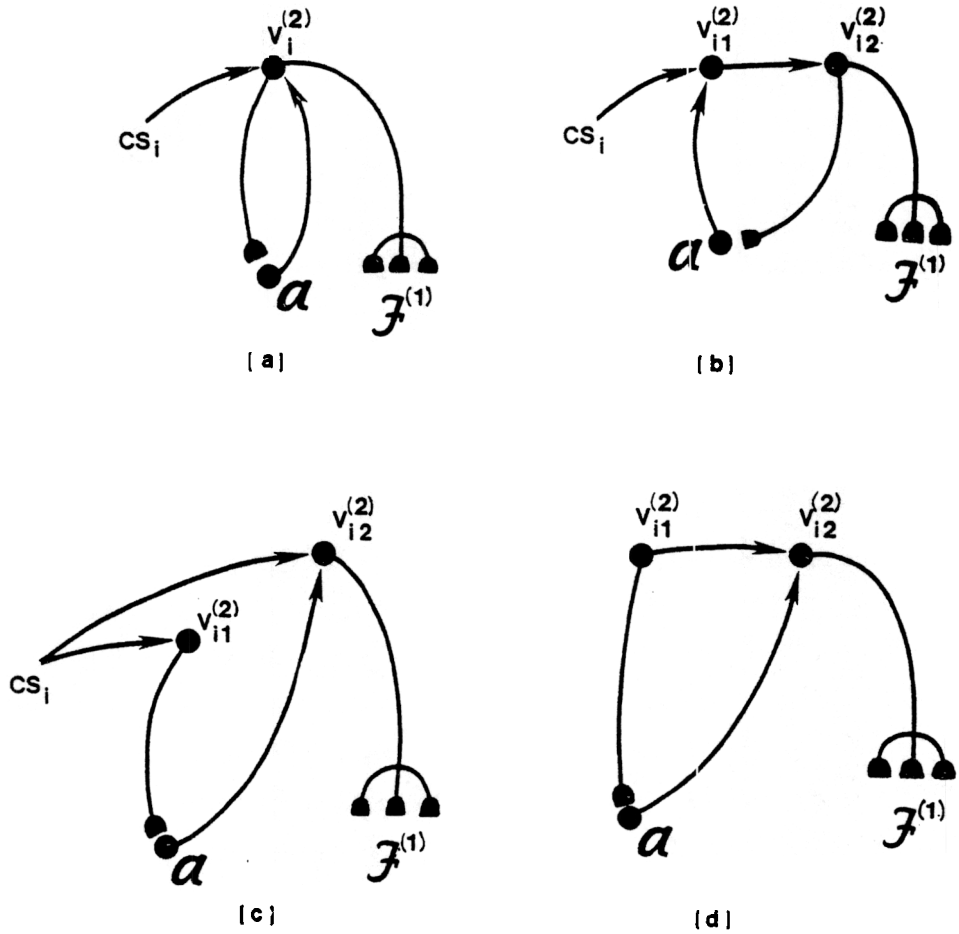

Fig. 6. Networks (*a*) and (*b*) do not work because their polyvalent cells can never fire. Network (*c*) can fire but it is incapable of storing the CS in STM. Network (*d*) is the minimal network capable of realizing postulates (A)–(D) and storing the CS in STM.

Stage $v_{i1}^{(2)}$ can only fire if it is aroused by $\mathscr{A}$, but $\mathscr{A}$ can only fire on recall trials if it is sampled by $v_{i2}^{(2)}$, which in turn can only fire if it is activated by $v_{i1}^{(2)}$. This is again an impossible arrangement. Consequently, two processing stages only help if they are connected in the correct fashion.

Both the networks in Figure 6*c* and 6*d* are possible, but Figure 6*d* enjoys an important advantage. In Figure 6*c*, the CS activates both stages $v_{i1}^{(2)}$ and $v_{i2}^{(2)}$. Stage $v_{i1}^{(2)}$ samples the arousal cells $\mathscr{A}$, and stage $v_{i2}^{(2)}$ can fire if it receives the CS input plus feedback from $\mathscr{A}$. Thereupon $v_{i2}^{(2)}$ samples the pattern at $\mathscr{F}^{(1)}$. The major disadvantage of this network is that sampling of $\mathscr{F}^{(1)}$ becomes impossible as soon as the CS shuts off. In Figure 6*d*, the CS activates stage $v_{i1}^{(2)}$, which thereupon sends a signal to stage $v_{i2}^{(2)}$ and samples $\mathscr{A}$. Stage $v_{i2}^{(2)}$ can fire when it receives the signal from $v_{i1}^{(1)}$ plus a feedback signal from $\mathscr{A}$. Thereupon $v_{i1}^{(2)}$ samples the pattern across $\mathscr{F}^{(1)}$. This network can be modified so that learning is still possible after the CS terminates.

To show what I have in mind, let me anticipate the argument a little by making the following observation. In instrumental conditioning experiments, the learning subject scans stimuli and emits behaviors before a reward or punishment occurs. What keeps the internal representations of the stimuli and behavior active after they terminate so that the later reinforcements can influence their interrelationships? In other words, at what stage of network processing does storage in short term memory (STM) occur?

There is a simple answer in Figure 6*d*: let the stage $v_{i1}^{(2)}$ reverberate in response to a CS signal (Figure 7). Then $v_{i1}^{(2)}$ can send persistent signals to stage $v_{i2}^{(2)}$. Both

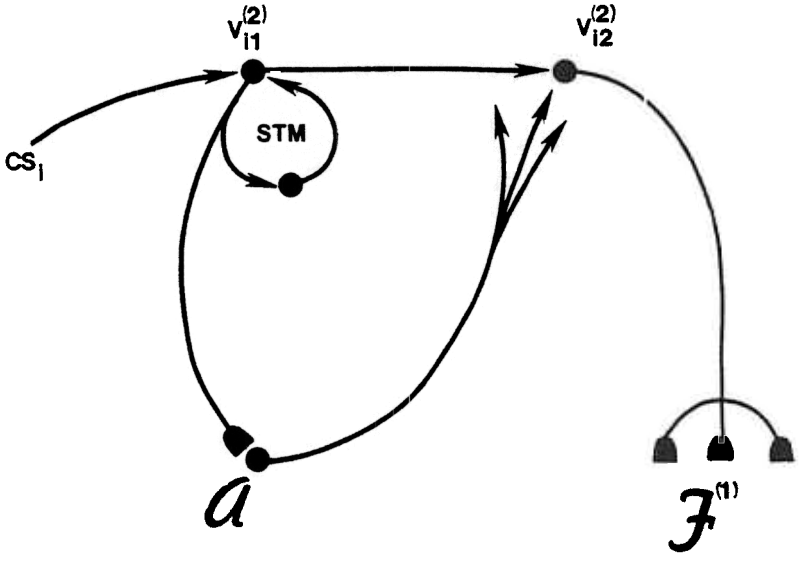

Fig. 7. CS_i can reverberate in STM at $v_{i1}^{(2)}$ as it emits sampling signals to $\mathscr{A}$ without being able to fire the polyvalent sampling cells $v_{i2}^{(2)}$.

stages 'code' their CS in the sense that they are selectively activated by it. However, stage $v_{i2}^{(2)}$ can only sample $\mathscr{F}^{(1)}$ if it also receives arousal from $\mathscr{A}$. Stage $v_{i1}^{(2)}$ persistently samples $\mathscr{A}$, but cannot learn to activate $\mathscr{A}$ until a UCS occurs. Thus STM reverberation can occur at $v_{i1}^{(2)}$ without erroneously eliciting sampling signals from $v_{i2}^{(2)}$.

This is not possible in the network of Figure 7*c*. If a reverberatory loop is added to $v_{i1}^{(2)}$, then nothing is accomplished, since $v_{i1}^{(2)}$ cannot activate $v_{i2}^{(2)}$ directly, and can only activate $v_{i2}^{(2)}$ indirectly after learning has already occurred, which prevents habit learning by $v_{i2}^{(2)}$ after its CS terminates. If a reverberatory loop is added to $v_{i2}^{(2)}$, and if $v_{i2}^{(2)}$ can fire this loop after the CS terminates and before learning occurs, then $v_{i2}^{(2)}$ can fire without arousal from $\mathscr{A}$, which is impossible. Hence Figure 6*d* is the only minimal anatomy that can solve the synchronization problem and can also accommodate CS-activated STM. Consequently, this anatomy will be used in the following discussion.

10. Secondary conditioning: a principle of equivalence for the anatomy of CS and UCS

Given that a CS can acquire UCS properties due to practice, we can conclude that important features of the anatomical representations of CS and UCS are often the same. To see this, let $\mathscr{S}$ be exposed to a sequence $E_1, E_2, \ldots$ of classical conditioning experiments. Denote the CS of E_i by CS_i and the UCS of E_i by UCS_i. Let the CS of E_i be the UCS of E_{i+1}; that is

$$\mathrm{UCS}_{i+1} = \mathrm{CS}_i, \qquad i \geqslant 1. \tag{2}$$

In other words every CS can become a future UCS and every UCS can have been a past CS except possibly UCS_1, on whose arousal properties the entire sequence of higher-order conditioning is built up.

The time scale of each conditioning experiment is short, on the order of minutes. I assume that new intercellular pathways cannot be wholly created or destroyed during the short time needed to go from any stage E_i to the next stage E_{i+1}. It follows that there exists a common anatomical representation for CS and UCS processing except possibly for UCS_1. By Section 9 every CS_i has a representation with at least two successive stages $(v_{i1}^{(2)}, v_{i2}^{(2)})$ of processing. Thus every UCS_i, $i > 1$, has the same representation. In other words, all CS and UCS inputs, except possibly UCS_1, are delivered to $\mathscr{F}^{(2)}$. Figure 8 illustrates the equivalence of these representations.

I call the property of common representation for CS and UCS pathways CS–UCS *path equivalence*. Path equivalence is the anatomical substrate that makes secondary conditioning possible.

Let us now summarize how the network in Figure 8 works during a classical conditioning experiment. Let the CS activate $v_{11}^{(2)}$ and the UCS activate $v_{21}^{(2)}$. When the CS occurs, $v_{11}^{(2)}$ fires and sends signals to $v_{12}^{(2)}$ and to the arousal population $\mathscr{A}$. Nothing else happens until the UCS arrives at $v_{21}^{(2)}$. This is because $v_{12}^{(2)}$ can only

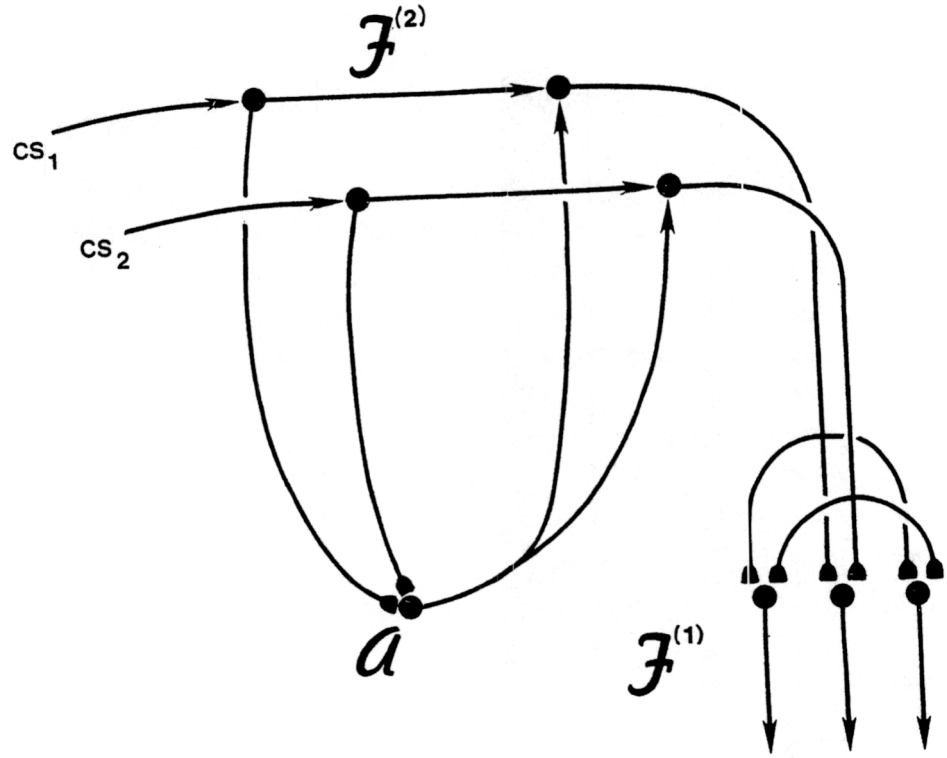

Fig. 8. Since each CS can rapidly acquire UCS properties, both CS's and UCS's possess a common anatomical substrate, except perhaps the primal, or pre-wired, UCS's.

fire if it receives an input from $v_{11}^{(2)}$ *and* from $\mathscr{A}$, but the signal from $v_{11}^{(2)}$ to $\mathscr{A}$ is initially too small to fire $\mathscr{A}$. When the UCS perturbs $v_{21}^{(2)}$, $v_{21}^{(2)}$ sends a signal to $v_{22}^{(2)}$ and to $\mathscr{A}$. The $v_{21}^{(2)}$ signals are large enough to fire $\mathscr{A}$, *because* the cue firing $v_{21}^{(2)}$ is a UCS. When $\mathscr{A}$ fires, it releases nonspecific signals to all cells $v_{12}^{(2)}$, $v_{22}^{(2)}$, $v_{32}^{(2)}$, . . . in $\mathscr{F}^{(2)}$. Now three things happen. First, since $v_{11}^{(2)}$ and $\mathscr{A}$ are both active, the LTM traces in the synaptic knobs of $v_{11}^{(2)}$ axons get stronger. When these traces get strong enough, the CS alone will be able to fire $v_{12}^{(2)}$. Second, the arousal signal from $\mathscr{A}$ combines with the UCS derived signal from $v_{21}^{(2)}$ at $v_{22}^{(2)}$, thereby firing signals from $v_{22}^{(2)}$ to $\mathscr{F}^{(1)}$. These signals elicit the UCS pattern in the populations of $\mathscr{F}^{(1)}$. Third, because the arousal signal from $\mathscr{A}$ is nonspecific, it also combines with the CS-derived signal from $v_{11}^{(2)}$ at $v_{12}^{(2)}$, thereby firing signals from $v_{12}^{(2)}$ to $\mathscr{F}^{(1)}$. These signals sample the UCS-elicited pattern at $\mathscr{F}^{(1)}$. Consequently, the CS begins to acquire UCS properties, both by learning to control the arousal pathway $\mathscr{A}$, and by learning to elicit (a component of) the UCS-induced pattern at $\mathscr{F}^{(1)}$.

Path equivalence also provides an elegant answer to the question: how does the UCS arouse the CS with just the right time lag to sample UCS onset at $\mathscr{F}^{(1)}$?

The same arousal which allows $v_{22}^{(2)}$ to read out the UCS pattern across $\mathscr{F}^{(1)}$ also allows $v_{12}^{(2)}$ to emit sampling signals to read in this UCS pattern.

11. Are drives energizers or sources of information?

Path equivalence has an important intuitive meaning. Consider the stages E_1 and E_2 of Section 8 for definiteness. What intuitive fact has changed when the equation

(3) $$CS_1 = \text{visual presentation of food}$$

is replaced by

(4) $$UCS_2 = \text{visual presentation of food}?$$

Visual presentation of food has taken on the significance of food by being conditioned to the arousal cells $\mathscr{A}$. The cue has acquired an internal meaning for $\mathscr{S}$. Arousal prepares $\mathscr{S}$ to be able to learn that the cue CS_1 signals forthcoming satisfaction of the internal demand for food. In particular, arousal is not merely an energizer, as Hull (1943) or Hebb (1955) suggested. It can have a cue function also, albeit a cue function concerning the internal state of $\mathscr{S}$ rather than the external state of the world.

Equations (3) and (4) describe cues that are related to hunger. Similar equations can also be written for cues that are related to thirst, sexual arousal, or fear, among other internal organismic states. Moreover, the same cue could have served as the CS for UCS's that are relevant to any of these states, just as a bell associated with shock can elicit autonomic signs of fear, whereas the same bell associated instead with food can elicit salivation. This observation is summarized in the postulate:

E. A given cue can be associated with any of several organismic states.

12. External facts versus internal demands

Postulate (E) implies that the arousal cells $\mathscr{A}$ are broken into several functionally distinct, but not necessarily independent, populations. These populations will be denoted individually by $\mathscr{D}_1, \mathscr{D}_2, \ldots, \mathscr{D}_m$, and collectively by $\mathscr{D}$, because they play the role of drive representations in the network. To represent hunger and thirst by different drive representations does not imply that the two representations are built up from disjoint cells. If the two representations do share cells, then every input to one representation will also deliver an input to the other representation whose relative intensity depends on the overlap of cells between the two representations.

Given the existence of several drive representations $\mathscr{D}_1, \mathscr{D}_2, \ldots, \mathscr{D}_m$, the previous discussion implies that each population $v_{i1}^{(2)}$ in $\mathscr{F}^{(2)}$ can sample several drive representations $\mathscr{D}_j$, and that each $\mathscr{D}_j$ sends signals to several populations $v_{i2}^{(2)}$. In just the same sense that signals from $\mathscr{D}$ to $\mathscr{F}^{(2)}$ are nonspecific, also signals

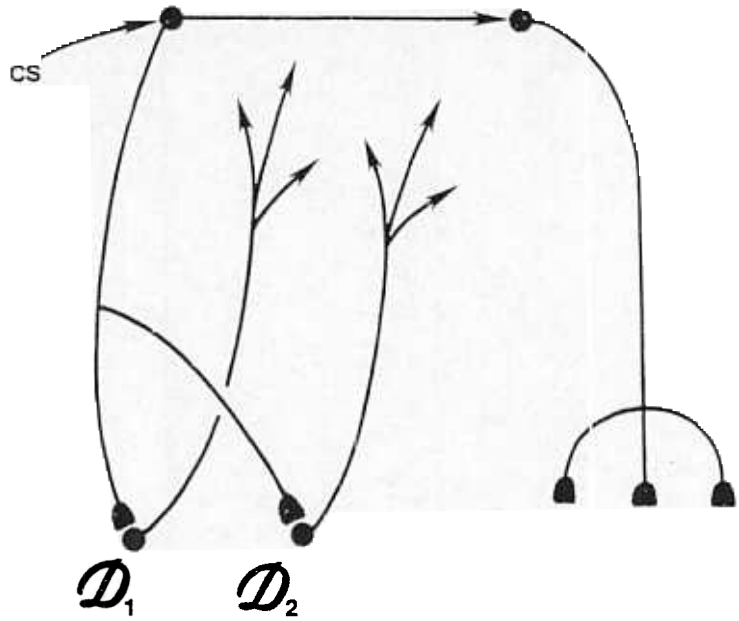

Fig. 9. Each CS can sample the several drive representations $\mathscr{D}_1, \mathscr{D}_2, \ldots, \mathscr{D}_m$. Each drive representation can, in turn, deliver nonspecific arousing signals to the CS's.

from $\mathscr{F}^{(2)}$ to $\mathscr{D}$ are nonspecific (Figure 9). Nonetheless, there are quantitatively more cells in $\mathscr{F}^{(2)}$ to which $\mathscr{D}$ projects than conversely. Since each $v_{i1}^{(2)}$ can now send signals to several $\mathscr{D}_j$'s, it is the source of an outstar. Consequently, each sensory representation can learn a spatial pattern of activity across the several drive representations as it is paired through time with UCS inputs. At each time, this spatial pattern summarizes the entire past history of drive state activations that occurred while its source cell was active.

An important issue concerns the *reciprocity* of connections between sensory and drive representations. If $v_{i1}^{(2)}$ projects to $\mathscr{D}_j$, does $\mathscr{D}_j$ always project to $v_{i2}^{(2)}$? If not, then $v_{i1}^{(2)}$ can be conditioned to drive $\mathscr{D}_j$ without $\mathscr{D}_j$ being able to release sampling signals from $v_{i2}^{(2)}$. Henceforth we always assume reciprocity for definiteness, although obvious modifications of our argument extend to the nonreciprocal case. We do not, however, assume *equipotentiality* of connections; namely, it is not necessary to our argument that each sensory representation projects to all drive representations, or that it projects with the same path strength to any pair of drive representations. Such asymmetries can influence if and how long learning can take, or even whether or not a particular learned behavior will be masked by a competing and more salient behavior. However they do not influence the primary network structure that our postulates imply. They are rather species-specific variations on this primary structure (Bitterman and Mackintosh, 1969; Seligman and Hager, 1972).

13. Internal facts versus external demands: existence of homeostatic, or drive, inputs

Thus far if a sensory representation in $\mathscr{F}^{(2)}$ becomes associated with a drive representation and a pattern across $\mathscr{F}^{(1)}$, it will release the $\mathscr{F}^{(1)}$ pattern whenever the sensory representation is activated. For example, a visual cue of food could always elicit eating behavior no matter how satiated the organism was. This property could, of course, have disastrous consequences. Speaking intuitively, the problem is clear: Some index of organismic need must exist to prevent unappropriate consummation. More generally, indices of the organism's internal states must be introduced to modulate its reactions to available external cues.

At this point in the theory, it could have happened that no plausible modification of the previously derived network dynamics could overcome this difficulty. Quite to the contrary, however, there exists a discernible symmetry in the networks, but a missing mechanism mars this symmetry, and its introduction will overcome our difficulty.

The symmetry is based on the fact that the representations of external cues project nonspecifically to the representations of internal drive states, *and conversely*. In a clear sense, the drive states are also representations of cues, albeit cues that are internal to the organism, rather than cues in the external world. The symmetry is marred by the fact that external cues representations can fire only if a specific external cue signal summates with a nonspecific internal drive signal. By contrast, the internal drive representations can fire whenever they receive a nonspecific external cue signal (Figure 10). To remove this asymmetry, a specific internal drive signal should also be necessary to fire the internal drive representation. If an external cue representation could fire without an external cue being

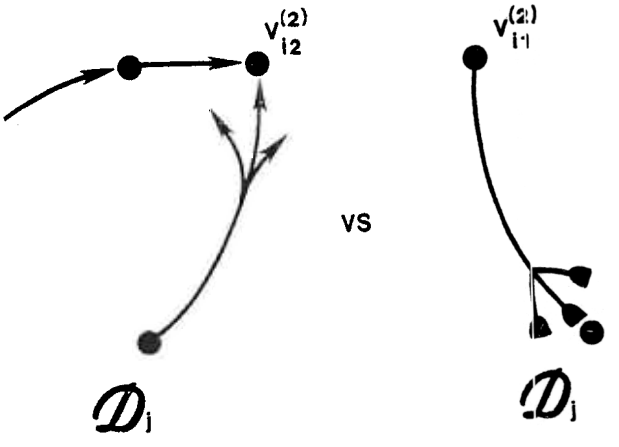

Fig. 10. A sensory representation $v_{i2}^{(2)}$ can fire only if it receives a specific external cue input plus a nonspecific internal drive input. By contrast, a drive representation $\mathscr{D}_i$ can fire whenever it receives a nonspecific external cue input. This asymmetry has unpleasant behavioral implications.

present, we would say that a type of hallucinatory event had occurred. Letting drive representations fire in the absence of specific internal drive signals is like permitting drive hallucinations to occur. Introducing specific drive inputs eliminates this network asymmetry and begins to overcome the satiety problem.

The above discussion can be organized around the following postulate.

F. Observable behavior is influenced by an organism's drives.

If this postulate were imposed in a theoretical vacuum, it would be just a vaguely stated triviality. Given the mechanisms that we have already derived, its minimal realization is to suppose, first, that there exist specific internal drive inputs to the drive representations whose sizes provide a measure of internal organismic needs; and second, that the drive representations can fire only if a sufficiently large specific internal drive signal occurs simultaneously with a sufficiently large nonspecific external cue signal (Figure 11). The drive representations are therefore also constructed from polyvalent cells.

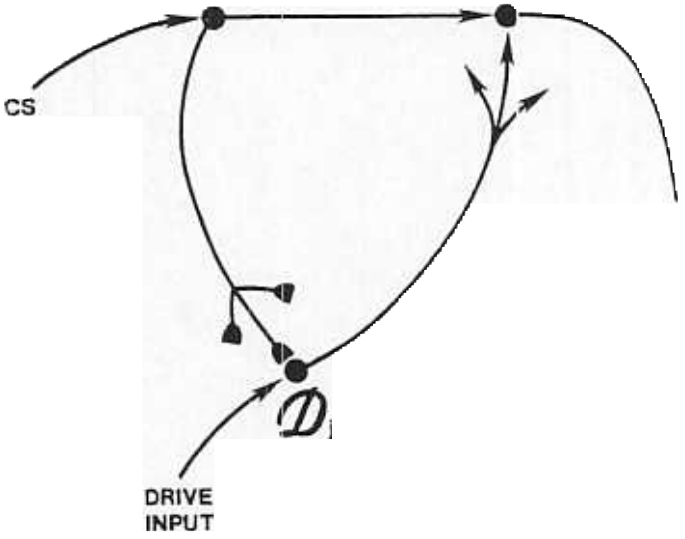

Fig. 11. The drive representations are also constructed from polyvalent cells whose firing requires convergence of a specific internal drive input plus a nonspecific external cue input.

The network's symmetry suggests the following question. Are both types of polyvalent cells anatomically homologous *in vivo*? I suggest that both cell types be identified with pyramidal cells. In particular, the $\mathcal{F}^{(2)}$ polyvalent cells are suggested to play the role of cortical pyramidal cells, whereas the $\mathcal{D}$ polyvalent cells are associated with hippocampal pyramidal cells (Shepherd, 1974). The feedback loop between $\mathcal{F}^{(2)}$ and $\mathcal{D}$ then becomes a rudimentary analog of the feedback exchange that takes place between cortex and hippocampus to regulate motivated behavior (Gabriel *et al.*, 1980; Grossberg, 1975).

14. Conditioned reinforcers, drives, incentive motivation, and habits

At this point, I can begin to psychologically interpret network mechanisms, use the interpretation to sharpen and modify classical psychological theories, and explain various data as manifestations of the adaptive designs that organisms have evolved to solve the synchronization problem.

Consider Figure 12. This figure labels network mechanisms in suggestive psychological jargon. *Reinforcement* acts by changing the spatial pattern that is coded in

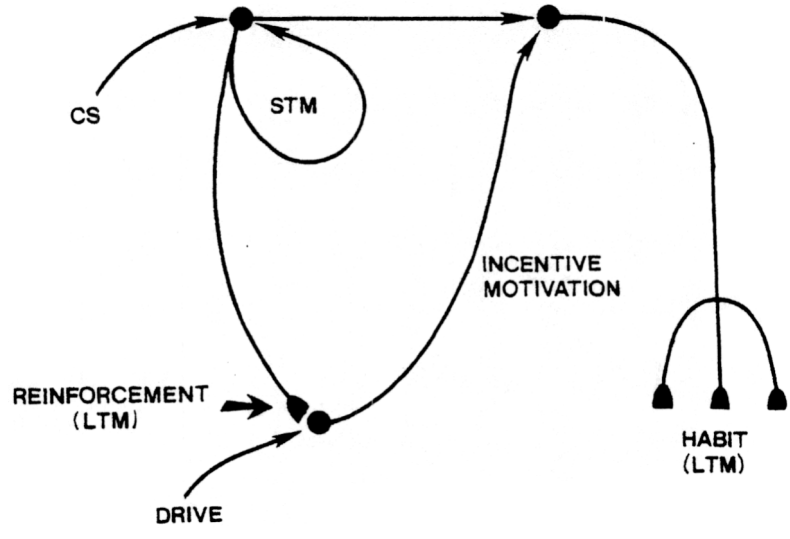

Fig. 12. The processing stages forced by the synchronization problem admit psychological labels which enable us to sharpen and modify classical psychological theories.

LTM when an active external cue representation samples the pattern of activity playing across the internal drive representations. As an external cue representation builds up large LTM traces in some of its pathways to the drive representations, its cue acquires *conditioned reinforcer* properties. In particular, its cue can be used as a UCS in a classical conditioning experiment. An organism's *drive* state is represented by the spatial pattern of drive inputs across its drive representations at a given time. Conditioned reinforcer and drive inputs merge at the drive representations, so that the LTM patterns that are learned by the conditioned reinforcer pathways are mixtures of reinforcement and drive information, rather than solely of reinforcement information.

When conditioned reinforcer and drive inputs combine with sufficient vigor at a drive representation, nonspecific arousal can be released. Each drive representation is a distinct source of arousal. Its arousal level is called the *incentive motivation* associated with the drive. Thus drive, reinforcer, and incentive signals are concep-

tually distinct. This distinction sharpens the familiar observation that motivation can be absent even though drive is high if an appropriate cue is missing. Also motivation can be absent even though an appropriate cue is present if the drive is low.

Finally, the incentive motivation joins with an external cue input to learn or read out an LTM pattern that represents *habit strength.* Read-out of such an LTM pattern can activate a command that controls a behavioral act. LTM changes thus occur at two loci in the network: they record the net pattern of reinforcer and drive data at the internal drive representations, and they record the pattern of behavioral commands at the habit representations.

15. Comparison with Hullian concepts

In Hull's theory, a drive energizes any and all behavior. As Bolles (1967, p. 183) clearly summarizes, according to Hull, 'only that behavior occurs which has the strongest associative connections, and drive merely determines the strength of the dominant response'. In the present theory, drive also has an energizing effect, since without it, nonspecific incentive signals cannot be released. However, drive also plays an informative and an associative role: informative because a drive input energizes only certain incentive pathways at the expense of others; associative because drive data can change the LTM patterns that are encoded by conditioned reinforcers.

Hull also suggests, at least formally, that drives and incentives play a symmetric role in their influence on habits. The Hullian law

$$(5) \qquad {}_{S}E_{R} = {}_{S}H_{R} \times D \times V \times K$$

says that drive level D, stimulus intensity V, and incentive motivation K all multiply habit strength ${}_{S}H_{R}$ to determine reaction potential ${}_{S}E_{R}$. In some ways, our networks support the Hullian formulation, since external cues, drives, and incentives all collaborate to read out the LTM patterns that encode network habits. However, in the network, external cues (akin to V) are gated by reinforced LTM patterns before they supply conditioned reinforcer inputs to the drive representations. Conditioned reinforcers and drive inputs (akin to D) determine incentive motivation (akin to K). Then incentive motivation and external cue inputs (akin to V) are gated by habits (akin to ${}_{S}H_{R}$). Let us adopt a Hullian type of notation ${}_{S}C_{R}$ for conditioned reinforcer LTM patterns that gate external cue signals on their way to drive representations. Then a Hullian analog of the network equations can be written as

$$(6) \qquad K = {}_{S}C_{R} \times V \times D$$

and

$$(7) \qquad {}_{S}E_{R} = {}_{S}H_{R} \times V \times K.$$

The most important features of (6) and (7) are that stimulus intensity V influences both K and ${}_{S}E_{R}$, that K is not independent of V and D, and that ${}_{S}C_{R}$ does not

equal ${}_sH_R$. However, even this refinement of the Hullian formalism omits most of the network's spatiotemporal structure. It is meant to celebrate Hull's intuition and to mark a path that leads beyond his formalism.

In the next few sections, I will use network mechanisms to analyze classical data, concepts, and theories. These classical contributions need to be mechanistically classified before later contributions that build on their shoulders can be differentiated into variations on old themes versus really new insights.

16. Data on conditioned reinforcers and drives

In the network, external stimuli become conditioned reinforcers when their pathways to drive representations are classically conditioned. Kelleher (1966, pp. 179–81) reviews experimental evidence that 'stimuli become conditioned reinforcers through respondent conditioning'.

In the network, a larger drive input can facilitate both performance and learning, albeit on different time scales. Performance is rapidly facilitated because a larger drive input can more vigorously activate its drive representation, which elicits incentive motivational signals, which energizes the performance of motivationally compatible habits. Learning is more slowly affected along both direct and indirect pathways. A large drive input can more vigorously activate its drive representation, which can then be directly sampled by the conditioned reinforcer LTM traces of active external cue representations. An indirect effect on learning occurs when the large incentive motivational signals elicited by the drive representation change the habits that can be sampled by active external cue representations. In various animals, weight loss is a good indicator of their motivation to learn a task leading to food reward (Stolurow, 1951; Bolles, 1967, Chapter 7). These data suggest that decrements in weight may cause proportional increments in the size of the drive input to the hunger drive representation. Such a drive input increment can energize both learning and performance in the manner suggested above.

Due to the interaction between performance and learning processes, the energizing effects of a drive input are not sufficient to explain network dynamics. Associative factors modulate a drive's efficacy, since the cues of an unfamiliar situation must become conditioned reinforcers and conditioned habit strength sources before they can efficiently control a learned behavior. In fact, when a hungry animal is introduced into an unfamiliar situation, its initial feeding behavior is often less vigorous than its feeding behavior later on, after it has already eaten enough to partially reduce its hunger. Such data, controlled for various alternative interpretations, have led to the conclusion that an animal's familiarity with the eating situation is a significant associative factor that influences the vigor of its feeding behavior (Bolles, 1967, Chapter 8; Moll, 1964).

17. Data on reinforcement

In the network, reinforcing effects occur when a large conditioned reinforcer input interacts with a large drive input to fire a drive representation, which is then sampled by an active cue representation. Bindra (1968) provides experimental

evidence that confirms this view of reinforcement. Bindra argues that there must exist a common neural locus where sensory inputs arising from incentive objects interact with the neural changes due to drive manipulation. These neural loci are the drive representations in the network. Scott and Pfaffman (1967) come to a similar conclusion from their studies of the hypothalamus, and Valenstein (1969) reviews more data of this type. Valenstein *et al.* (1970) report data in which 'hypothalamic stimulation . . . seems to create conditions which excite the neural substrate underlying well-established response patterns. . . . Discharging this sensitized or excited substrate is reinforcing and it can provide the motivation to engage in instrumental behavior . . . rats which display stimulus-bound eating prefer the combination of food and brain stimulation to brain stimulation. . . . The brain stimulation does not fully activate all the neural circuits underlying reinforcement . . .' These data are explicable if we assume that network drive representations include, or are activated by, these hypothalamic sites, and that brain stimulation acts like a drive input, albeit an artificial one. Then the network interpretation of the Valenstein *et al.* data is that conditioned reinforcer inputs must be bolstered by drive inputs in order to activate the drive representations, which thereupon simultaneously release incentive signals and cause learned changes in the conditioned reinforcer LTM patterns.

18. Data on self-stimulation

Does electrode current in the hypothalamus act like a drive input as the Valenstein *et al.* data suggest? Olds (1955, 1958, 1977) reviews data on self-stimulation behavior that support this view. Olds showed that a rat will learn to push a lever at high rates if it activates an electrode placed at suitable sites in its lateral hypothalamus. Often the rat will turn away from food or a mate to press the lever with great vigor until it becomes exhausted. Sites which elicit such consummatory self-stimulation behavior have been generically called the pleasure center. More detailed studies of self-stimulation suggest that the electrode input acts like a specific drive input. An electrode placed in an area associated with hunger loses its reinforcing effect when the animal is satiated, and the rate of lever pressing for self-stimulation increases when the animal is hungry. These data suggest that the usual hunger drive input summates with the artificial electrode input. In a similar fashion, an electrode placed at certain other hypothalamic sites will elicit faster lever pressing after androgen is injected. Androgen normally has the effect of motivating sexual behavior. At such loci, self-stimulation disappears almost completely after castration, which presumably eliminates the usual drive input to these loci. Similarly, higher current levels are needed to achieve self-stimulation as the androgen level subsides, indicating once again the energizing effect of drive inputs. This latter effect can be reversed by injections of testosterone propionate in oil.

By contrast, if androgen is injected when the electrode is located at hunger-related sites, then the lever-pressing rate decreases. If the animal is deprived of food when the electrode is at sex-related sites, then the lever-pressing rate again

decreases. Thus the various drive loci reciprocally inhibit, or compete with each other, whereas electrode input at a given drive locus works synergistically with the drive input at this locus. Section 28 discusses how drive competition is realized by network mechanisms.

In the network, self-stimulation can be explained as follows. At the moment when the exploring animal accidentally presses the lever, sensory cues of events immediately preceding this act are active in sensory STM, and the motor commands for pressing the lever are active in motor STM. The lever press releases a large artificial drive input to its drive representation, large enough in fact to fire the drive representation. The conditioned reinforcer pathways of the active cues which abut the drive representation are hereby conditioned. Also incentive motivation is released that enables the active cues to learn the motor commands controlling lever-pressing behavior. As trials proceed, the sensory cues of the lever gradually gain control over lever-pressing behavior, because their conditioned reinforcer and conditioned habit pathways are progressively strengthened on each trial by the electrode current that reliably follows the lever press.

This description helps us to understand why self-stimulation behavior is labile. In particular, self-stimulation shows poor resistance to extinction, poor carry-over of performance between learning sessions, and is a poor source of secondary reinforcement (Bolles, 1967, Chapter 9; Mogenson *et al.*, 1965; Stein, 1958). A basic difficulty is that, without electrode current available to continually reinforce the drive and habit pathways, the shifting pattern of drive and conditioned reinforcer inputs can act to remove adequate input for firing the drive representation or to competitively inhibit the drive representation that had received electrode input. Other sources of lability are disconfirmation of expectancies after continuously rewarded trials and the lack of experimental contingencies between the animal's expectancies and the onset of current (Grossberg, 1982*a*,*b*).

19. Reinforcement without drive reduction

Self-stimulation data were a major embarassment for Hull's central thesis that an event is reinforcing if it reduces a drive. Hull's idea seems to be supported by commonsensical experiences such as: you will learn a task to eat, and thereby reduce your hunger drive; or you will learn a task to escape shock, and thereby reduce your pain and fear. However, what drive is reduced when an animal pushes a lever to pump as much electric current as it can into its lateral hypothalamus? In a clear intuitive sense, the animal is working to increase, not to decrease, electric current. Moreover, why does an animal self-stimulate more if its hunger drive is increased and the electrode is placed in a hunger-related locus?

One might argue that the self-stimulation paradigm is so abnormal that the animal's behavior is not really reinforced, but only seems to be reinforced. However, once the floodgates were opened, many other behaviors could be cited that seem to violate the drive reduction principle. As an early example, Sheffield (1951) found that sexually naive male rats will learn an instrumental response if they are

rewarded by being allowed to copulate with receptive females, but not allowed to ejaculate. What drive is reduced in this situation?

We are now faced with an important dilemma. If drive reduction is not really the mechanism of reinforcement, then why does common sense so strongly suggest that it is? A deeper version of the dilemma is this. Is it correct to claim that shock reduction and hunger reduction are both examples of a common drive reduction mechanism? There is a clear sense in which shock is an aversive cue, so that reducing it might well be positively rewarding. In other words, reducing a negative cue can have a positively rewarding effect. But is hunger always aversive? Is sexual desire always aversive? What is aversive about anticipating a marvelous dinner and eating it with gusto? It would be aversive, to be sure, if the dinner were cancelled at the last minute. Is this aversive reaction due directly to hunger, or is it due to the frustration that is triggered when we learn that the expected feast has been called off?

If we agree that hunger or sexual desire can be positive drives, not negative drives like fear, then why should reducing them be reinforcing? The commonsense basis for believing in drive reduction hereby collapses. But then why do hunger and sexual desire sometimes seem aversive? Here we must distinguish between very high levels of these drives and normal levels. At normal levels, one can easily confuse the frustration caused by delayed gratification of the drive with the drive itself.

In the network, a positive drive input must be high before reinforcement can occur, because otherwise its drive representation cannot fire and cause LTM changes in abutting conditioned reinforcer pathways. The drive reduction that follows consummatory behavior is not, however, reinforcing. Instead, it prevents consummation after satiety occurs. This type of drive reduction occurs slowly in time. The sudden reduction of a reinforcing cue or the sudden nonoccurrence of an expected reinforcer can also have reinforcing effects, but these rapid events do not reduce a drive input, although they do modify the activity of the drive representations (Grossberg, 1982*a*,*b*).

20. Go mechanism, amplifiers, now print

Several authors have proposed alternatives to the drive reduction hypothesis to explain data about reinforcement. Each author developed his own vocabulary to describe his concepts, but all of them seem to have been building towards similar mechanisms. I will provide a comparative analysis to clarify some advantages of a network theory.

Miller (1963) introduced 'go' mechanisms that 'act to intensify ongoing responses to cues', that 'are subject to conditioning with contiguity being sufficient', such that 'the strength of the CR is determined to a great degree by the strength of the UCS', and 'when a chain of cues leads to a UCS for the "go mechanism", it is most strongly conditioned to those nearer to the UCS'.

Miller's mechanism is analogous to an incentive motivational signal. Incentive

motivation acts to 'intensify ongoing responses to cues' by controlling the size of signals in the habit strength pathways, and can be 'conditioned with contiguity being sufficient' in the conditioned reinforcer pathways. Moreover, CS strength is 'determined to a great degree by the strength of the UCS' due to two factors acting together. The UCS input at the drive representations directly enhances performance in response to the CS by eliciting incentive motivational signals that amplify the CS-activated signals in the habit strength pathways. The UCS input also indirectly strengthens the CS by enhancing conditioning of the CS-activated conditioned reinforcer pathways, and thereby enabling the CS to activate stronger incentive motivational signals on later trials. These two effects can be experimentally distinguished because the direct effect acts quickly whereas the indirect effect builds up slowly.

Miller's pioneering concepts provide a useful intuitive description, but one that is weakened by lumping together several mechanisms that are invisible without a real-time theory.

Estes (1969) develops analogous concepts within his framework of stimulus sampling theory. He suggests that the occurrence of a response requires summation of input from external stimulus and internal drive sources. Drives and rewards serve as response amplifiers. On learning a trial, the organism $\mathscr{S}$ draws a sample of available discriminative cues and scans these cues until an element is processed which is connected with a permissible response. This response will be evoked only if an amplifier element appropriate to the response is simultaneously scanned. Stimuli can be conditioned to amplifier elements by contiguity, and the base rate of amplifier elements associated with a given drive increases as $\mathscr{S}$'s need increases.

The amplifier elements of Estes' theory play the role of incentive motivational signals in the network. Many of the intuitive distinctions described by Estes' theory are also found in the network theory. The two theories nonetheless differ in important ways that have limited further development of the Estes theory, but not of the network theory. The concepts of the Estes theory are expressed in the probabilistic language of stimulus sampling theory. External cues and amplifier elements are said to be scanned, presumably by some probabilisitic serial mechanism, and conditioning of cues to amplifier elements changes the probability of a successful joint scan of cues and amplifiers. By contrast, in the network theory, the probabilities of scanning amplifier elements are replaced by activity patterns that exist across all the drive representations at each time. Thus a serial scan is replaced by a parallel pattern. The activities of drive representations are not probabilities, nor are they scanned. It is hard to overemphasize the importance of this distinction. A larger drive activity will, other things equal, cause a higher probability of that drive influencing observable behavior. However, on each trial the drive activity influences the computations that determine observable behavior. It is not possible to fail to scan an active drive representation on a given trial.

Estes' theory omits other distinctions that are important in the network theory. For example, Estes suggests that internal drives set the base rate of amplifier elements, and that external cues modulate this rate. The network analog of ampli-

fier elements is incentive motivational signals, or alternatively the activities at drive representations that induce these signals. However, internal drive inputs are conceptually distinct from incentive motivational signals in the network. It is just as inadmissible to let drive inputs fire incentive motivational signals—i.e., set the base rate of amplifier elements in the absence of conditioned reinforcer signals—as it is to let external cue signals sample habit strengths in the absence of incentive motivational signals.

In the network theory, patterned LTM traces vary slowly through time and gate rapidly fluctuating signals from the external cue representations to the drive representations. In the sampling model, these two distinct processes are lumped into the probability that a scanned cue will be associated with an amplifier element. A conditioned reinforcer LTM trace can, however, be very large even if, due to the momentary STM inactivity of the cue representation, there is a zero probability of 'scanning an amplifier element' at that time. Thus, the network's way of representing the effects of prior reinforcement do not neatly correspond to the probabilistic concepts of the sampling theory.

More generally, the framework of stimulus sampling theory cannot easily represent the internal geometry of specific and nonspecific pathways, or the several time scales on which STM and LTM traces fluctuate. These deficiencies become especially evident when the network theory incorporates antagonistic rebound mechanisms (Grossberg, 1972*c*, 1981*a*, 1982*a*,*b*) and expectancy mechanisms (Grossberg, 1976*b*, 1980*a*, 1982*a*,*b*) to understand various conditioning and attentional phenomena. In a stimulus sampling context, these concepts are hard to motivate or to represent, but in a network framework they arise in a natural fashion. Thus despite its great heuristic value as a tool for classifying a variety of learning experiments, the stimulus sampling theory becomes increasingly unwieldy and inaccurate as it attempts to represent the intervening variables that govern complex learning behavior.

Logan (1969) claims that rewards 'excite' rather than 'strengthen' habits by providing 'incentive motivation' that favors their execution. Though the distinction between 'exciting' and 'strengthening' a habit might seem obscure, with the network theory as a guide, a possible mechanistic interpretation is suggested. The reward elicits incentive motivational signals that allow the habit to be released. This is the 'exciting' effect of a reward. I claim, however, that the reward can also 'strengthen' the habit in two distinct ways, albeit indirectly and on a slower time scale. One strengthening effect of reward is due to the conditioned reinforcer learning that it can trigger. Such learning can strengthen the incentive motivational signal which a reinforcing cue can elicit to 'energize' the habit on later performance trials. A second 'strengthening' effect of reward is more direct. It is due to the fact that a larger incentive motivational signal can cause a larger sampling signal to be emitted from a cue's polyvalent cells to the habit representation. A larger sampling signal implies a faster rate of habit learning, which can 'strengthen' the habit measured after a fixed number of learning trials.

Livingston (1967) also has a similar mechanism in mind in his discussion of a 'Now Print' mechanism that can control the learning rate.

It is remarkable that so many languages have been used to describe the same mechanisms. The rigorous explication of these mechanisms will hopefully unify the languages used to discuss them.

21. Data on incentive motivation

A male animal left alone will busily do the many things characteristic of his species, such as grooming, eating, exploring. He does not look sexually motivated. However, if a female animal is presented to him, his behavior can dramatically change (Beach, 1956; Bolles, 1967, Chapter 7). This example distinguishes between sex drive and observable motivated sexual behavior in the presence of appropriate external stimuli. Drive without external cues need not elicit motivated behavior. In the network theory, drive inputs cannot fire drive representations unless they are supplemented by an auxiliary input source such as reinforced signals from cue representations.

Furthermore, external stimuli without drive need not elicit motivated behavior. Seward and Proctor (1960) and Seward *et al.* (1958, 1960) found that if an animal is not hungry, then no amount of food will be adequate to reinforce its behavior. In the network theory, without drive inputs, no amount of conditioned reinforcer input can fire the drive representations. These experiments support the hypothesis that the polyvalent cells of drive representations need simultaneously converging reinforcing and drive inputs to vigorously fire.

An analogous type of experiment shows that cues can be trained to elicit appropriate behavior even if two or more drives are simultaneously active. Kendler (1946) trained rats who were simultaneously hungry and thirsty in a T maze. Food was on one side of the T maze and water was on the other side. The animals were forced alternately to the two sides, and rapidly learned the discrimination. When they were just hungry or just thirsty on a test trial, the rats went to the correct side with high probability. As Bolles (1967, Chapter 9) notes, this plausible result is embarrassing to a simple stimulus–response theory of conditioning, since on training trials, turning left and turning right should be associated to both the hunger and the thirst cues, but then how is the discrimination learned? In the network theory, the explanation is simple. The drive inputs of hunger and thirst cannot by themselves fire their drive representations. When food is eaten, its reinforced path to the hunger drive representation combines with the hunger drive input to selectively fire the hunger drive representation. The external cues associated with the food thereupon become conditioned reinforcers with respect to the hunger drive representation. A similar argument holds when water is drunk. Variations on the original Kendler experiment supported a role for incentive motivation (Kendler and Levine, 1951; Kendler *et al.*, 1952), but provided no mechanistic description of how incentive motivation differs from drive and reinforcement. Some recent theories of conditioning have fallen into difficulties

analogous to those implied by the Kendler (1946) data, because they do not include the notions of a drive representation and incentive motivation (Hall and Pearce, 1980). Similar ideas to those described herein have been used to overcome these difficulties (Grossberg, 1982*b*).

In Hullian theory and its variations, incentive motivation is often denoted by r_G. Tracking the existence and properties of r_G has led to many beautiful experiments and ingenious, if sometimes tortured, conceptualizations. Bolles' (1967) book provides a stimulating account of many classical efforts. Often the mysterious r_G was invoked because some other concept, such as drive or reinforcement, was inadequate, as in the above experiments. One did not know what r_G itself was, but one could often say what the other concepts were not. For example, many experimentalists claimed that incentive motivation is mediated by an anticipatory goal reaction. As Bolles (1967, p. 332) wrote: 'Drives push and incentives pull'. Drives are unlearned, whereas incentives are learned. Drives described the organism's momentary state, whereas incentives summarize the organism's history. Incentivey are thus used to explain many of the performance variations that are not directls due to momentary drives or to associative properties of habits. For example, Crespi (1942, 1944) emphasized the motivational properties of incentives by running different groups of rats in a straight alley to large or small quantities of food. Not surprisingly, large-amount animals performed better than small-amount animals. Then Crespi switched half of each group after twenty trials from high to low or from low to high. Performance changed rapidly to new levels commensurate with the new reward contingencies. The performance shifts were too fast to be explained by changes in drive level or changes in associative habits. In the network, changing the level of reinforcement rapidly changes the size of inputs due to the food itself, and thus the ambient level of polyvalent cell activation, even without a change in drive input level, or a change in conditioned reinforcer properties of external cues.

The classical r_G concept faces several fundamental difficulties. One difficulty is summarized by Bolles' (1967) phrase: 'The anticipatory goal response is the only serious proposal that has been made for a mechanism to account for incentive motivation' (p. 336). The Mackintosh (1974) view that was summarized in Section 1 is a more recent version of this idea. The problem is that this viewpoint lumps together several distinct processes under the r_G label. Truly anticipatory behavior often involves a behavioral plan whose unfolding is regulated by expectations that are matched against environmental feedback (Grossberg, 1978*b*, 1980*a*). By contrast, incentive motivation *per se* can seem to guide anticipatory behavior without either plans or expectancies being operative. For example, suppose that external cues excite a generalization gradient of related cue representations, and that all of the excited cue representations sample the drive and habit representations with a strength that depends on how excited they get due to the external cue. To test whether the animal 'anticipated' a certain cue or outcome, one might use related cues versus unrelated cues on test trials. Or one might use related versus unrelated response measures. A differential effect on performance might suggest

the action of anticipatory goal responses, but really no cognitive type of expectancy or anticipation need be involved. There exist basic reasons for the operation of expectancies even in simple conditioning paradigms (Grossberg, 1981*a*, 1982*a*,*b*), but for now let us note that incentive motivation and expectancy matching are wholly distinct mechanisms that tend to get badly obscured in the r_G literature.

Some of the multiple roles ascribed to r_G are legitimate, as when Bolles (1967, p. 354) asks if incentives 'reinforce instrumental behavior, motivate it, or simply provide stimulus control for it?' In the network, the answer is: all of the above, but using different pathways that can operate on different time scales.

22. Secondary reinforcement and Hull's paradox

Hull's use of drive reduction together with r_G created a serious paradox (Bolles, 1967, p. 355). Hull claimed that a stimulus associated with drive reduction acquires secondary reinforcing power. It can act the same way as drive reduction does in later learning experiments. Hull supposed that the mechanism for this is r_G. Thus the occurrence of r_G should act like drive reduction. However, the motivating effects of a secondary reinforcer were also ascribed to an increase in incentive motivation, or r_G. How can r_G both be drive reducing and drive increasing? Given Hull's formulation, it seemed hard to argue that r_G is both reinforcing and motivating.

This problem is overcome in the network theory. Drive reduction is abandoned, and path equivalence shows how a UCS can be reinforcing, via its effects on the LTM of other active conditioned reinforcer pathways, and motivating, via its effects on the firing of habit strength pathways.

23. Late nonspecific potential shifts

Some seemingly paradoxical data have been collected by doing discrimination learning on animals with implanted electrodes. John and Morgades (1969) reported that, in trained animals, discriminative stimuli elicit characteristic responses that are distributed rather uniformly across extensive cellular regions, and that these uniform reactions manifest themselves as an increase in the late components of evoked brain potentials as training goes on. The paradox is as follows. Why should an increase in the animal's discriminative ability correspond to a more uniform distribution across space of cell potentials? Why don't the potentials get more discriminative also?

Several authors have interpreted this result by claiming that, as discrimination improves, the 'information' about the discrimination is spread uniformly across the network, akin to an equipotentiality or hologram type of concept. This idea creates several problems when we ask how nerves can retrieve this uniformly scattered information, since all discriminations will become uniformly distributed across the same neural tissue. In a holographic theory, such decoding relies on the existence of precisely calibrated periodic sampling probes. Within the active neuropile, a much more chaotic temporal behavior prevails.

In the network theory, this interpretation seems to be unnecessary at best. As discrimination training proceeds, the external cues gain control over incentive motivational pathways as they become conditioned reinforcers. Since the incentive motivational pathways are nonspecific, the cues can deliver their signature quite uniformly to all the sites that receive the incentive motivational signals. These uniform signals are 'late', because it takes longer for them to feed through the drive representations and then back via incentive motivational pathways than it does for the cues to directly activate sensory STM. However, the 'information' in the network is not carried by these signals. The information is carried in the polyvalent properties that determine which cells will fire, and in the LTM patterns that will subsequently be read out.

It is somewhat surprising that John himself did not reach this conclusion, since John (1966, 1967) has also reported that neocortical polyvalent cells require both CS and UCS input in order to fire. Perhaps John's oversight was due to the theoretical property that after discrimination has taken place, this is no longer true; then the CS alone is a sufficient external cue to fire its pathway. John's data thus contained a gap, and because he took the data at face value, he was led to a paradoxical hypothesis.

The contingent negative variation (CNV) is a slowly varying cortical potential shift that is a likely neural substrate of incentive motivational signals. Walter (1964) hypothesized that the CNV shifts the average baseline of the cortex by depolarizing the apical dendritical potentials of cortical pyramidal cells and thereby priming the cortex for action. This is why the CNV has been associated with an animal's expectancy, decision (Walter, 1964), motivation (Irwin *et al.*, 1966), preparatory set (Low *et al.*, 1966) and arousal (McAdam, 1969). The CNV has also been described as a conditionable wave. Thus far in the network theory, the incentive motivational pathway can be conditioned only indirectly when a cue acquires conditioned reinforcer properties that are reflected by that cue's growing ability to elicit incentive motivational signals. We will return to the question of whether incentives can be directly conditioned in Section 25.

24. An emergent neocortical analog

We have now amassed enough empirical evidence to suggest some anatomica analogs of the wiring diagram in Figure 13 (Grossberg, 1978*b*). These analogs are suggested tentatively for several reasons. Most importantly, when the synchronization postulates are supplemented by new postulates, more processing stages will be imposed on the evolving network. Consequently, although anatomical interpretations of some network processing stages can be identified with reasonable assurance, the stages just before and after these stages have a more ambiguous anatomical interpretation. Also species-specific variations on network themes must be anticipated. Anatomical markers are nonetheless useful to facilitate comparison with neural data, and also to indicate that the formal network stages have plausible neural interpretations. Below two possible interpretations of Figure 13 will be

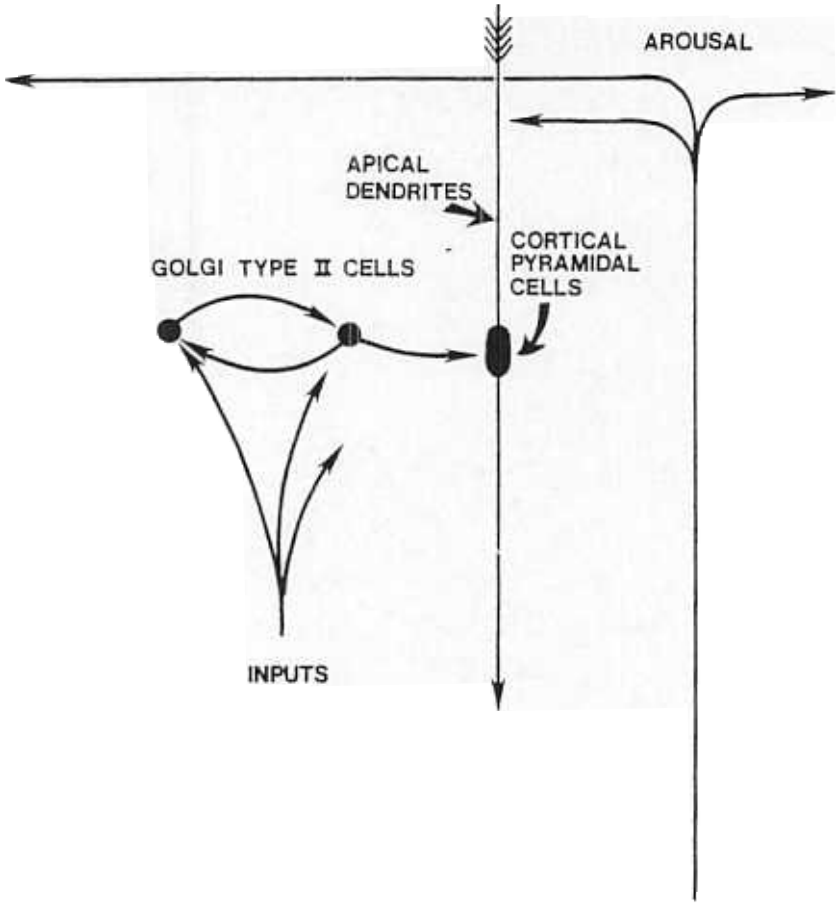

Fig. 13. An interpretation of network dynamics in terms of nonspecific arousal afferents to the apical dendrites of cortical pyramidal cells, and CS-stimulated reverberation of Golgi type-II interneurons whose output stimulates the pyramidal cells.

suggested. Each interpretation leads to definite questions about cortical dynamics.

Both interpretations depend on identifying the second stages $\{v_{i2}^{(2)}\}$ of sensory processing with cortical pyramidal cells, which are the output cells of cerebral neocortex (Shepherd, 1974; Sholl, 1956). The work of John (1966, 1967) on polyvalent cortical cells and of Walter (1964) on the CNV both suggest that we identify these polyvalent cells with cortical pyramidal cells. Given this interpretation, what are the cells subserving the other stages? The main difference between the two interpretations concerns the issue: Do the first sensory stages $\{v_{i1}^{(2)}\}$ occur in the cortex or not?

Suppose 'yes'. Then the cells which subserve STM storage at the first stage of the sensory representation are cortical interneurons. These interneurons are excited

by cortical inputs, mutually excite each other, and send excitatory signals to pyramidal cells (Figure 14*a*). We identify these interneurons with Golgi Type II cells (Crosby *et al.*, 1962; Peters *et al.*, 1979). The pyramidal cells cannot fire unless they also receive arousal inputs that, following Walter (1964), are assumed to prime their apical dendrites (Figure 14*b*). The main source of ambiguity con-

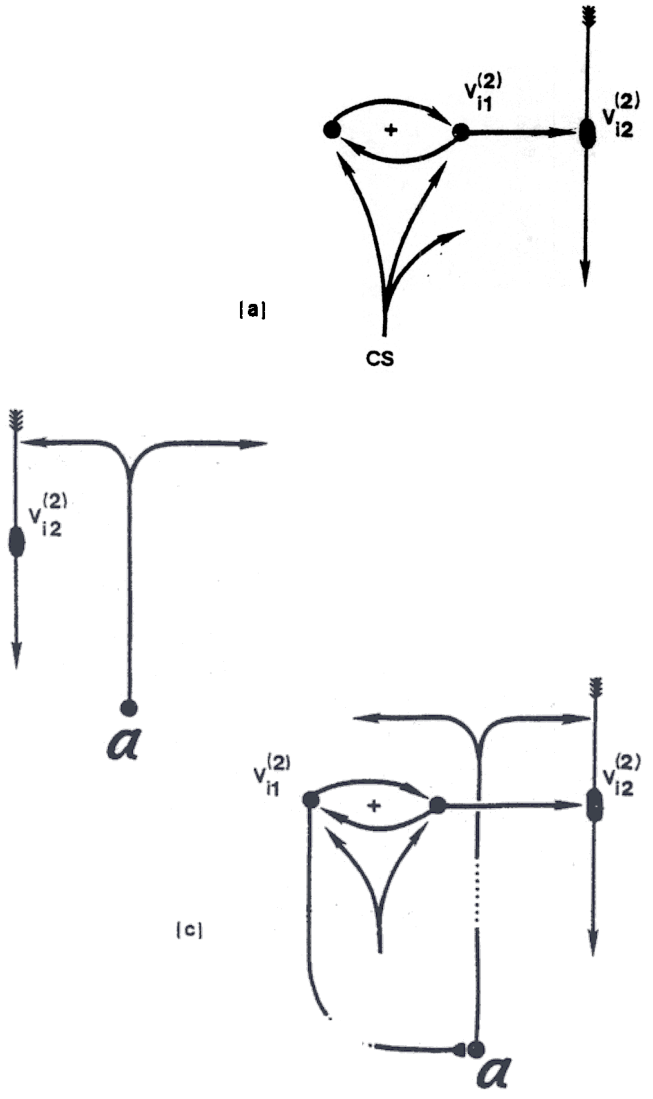

Fig. 14. (*a*) CS inputs excite Golgi-type II interneurons at $v_{i1}^{(2)}$ which reverberate and excite the pyramidal cells $v_{i2}^{(2)}$; (*b*) The pyramidal cells also receive nonspecific arousal inputs from the drive representations; (*c*) If the STM-reverberation is intracortical, the nonpolyvalent cortical output cells exist which can sample the drive representations.

cerns another class of cortical output cells that would also have to exist. These are the output cells in the first stages of the sensory representations, which can fire to the drive representations even without incentive motivational support (Figure 14*c*). If all cortical output cells are pyramidal cells (Shepherd, 1974), then these pyramidal cells would differ both in their firing rules and their output targets from the pyramidal cells depicted in Figure 14*a*. Their output targets would be drive representations that are in, or associated with, the hypothalamic and hippocampal sites at which drive inputs are evaluated, as the results of Scott and Pfaffman (1967) and Valenstein *et al.* (1970), as well as numerous other investigators suggest (Grossman, 1967; Haymaker *et al.*, 1969; Olds, 1977; Swaab and Schadé, 1974). These outputs could be triggered by cortical inputs, supplemented perhaps by STM reverberation, but without the support of incentive motivational signals. Perhaps these output cells are small pyramidal cells whose dendrites abut the cortical input pathways and interneurons with close proximity. Since many pyramidal cells have apical dendrites that rise to the upper layers of the cortex, this anatomical interpretation will fail if all pyramidal cells can be shown to require nonspecific incentive motivational input to their apical dendrites before they can fire.

Two main points in this interpretation are useful even if it does not survive the test just described. First, the division of a sensory representation into two successive stages is not a mysterious anatomical notion. It can, for example, be realized as the distinction between interneurons and output cells. Second, the polyvalent rules for output cell firing suggest one reason why certain neural tissues, such as neocortex and hippocampus, receive inputs which are segregated into distinct lamina. By running the dendrites of output cells through these lamina, one can control which combination of inputs is needed to fire the cells. This is especially true if the output cells with larger cell bodies and axons also have larger dendritic trees. Then the output cells that can fire most directly to spinal centers will require convergence from more input sources, and thus a less ambiguous configuration of input data, to fire their larger cell bodies (Grossberg, 1978*b*).

The alternative anatomical interpretation is suggested by the possibility that cortical pyramids cannot be fired without incentive inputs to their apical dendrites. In species where this is true, the first stage of sensory representation would be represented at an anatomical level prior to the cortex, such as within a specific thalamic nucleus (Anderson and Eccles, 1953; Crosby *et al.*, 1962; Gabriel *et al.*, 1980; Grossman, 1967; Macchi and Rinvik, 1976; Tsumoto *et al.*, 1978). Then the first sensory representation and its STM reverberation would exist in the thalamus, from which ascending cortical projections and descending limbic projections (Figure 15*a*) would arise. The drive representations would again deliver input to apical dendrites of pyramidal cells, but now the problem of firing cortical outputs without simultaneous incentive motivational inputs is averted. Also, cortical interneurons could reverberate among themselves and feed their activity to pyramidal cells (Figure 15*b*), but these interneurons would no longer fire the first sensory representations, which no longer exist in the cortex.

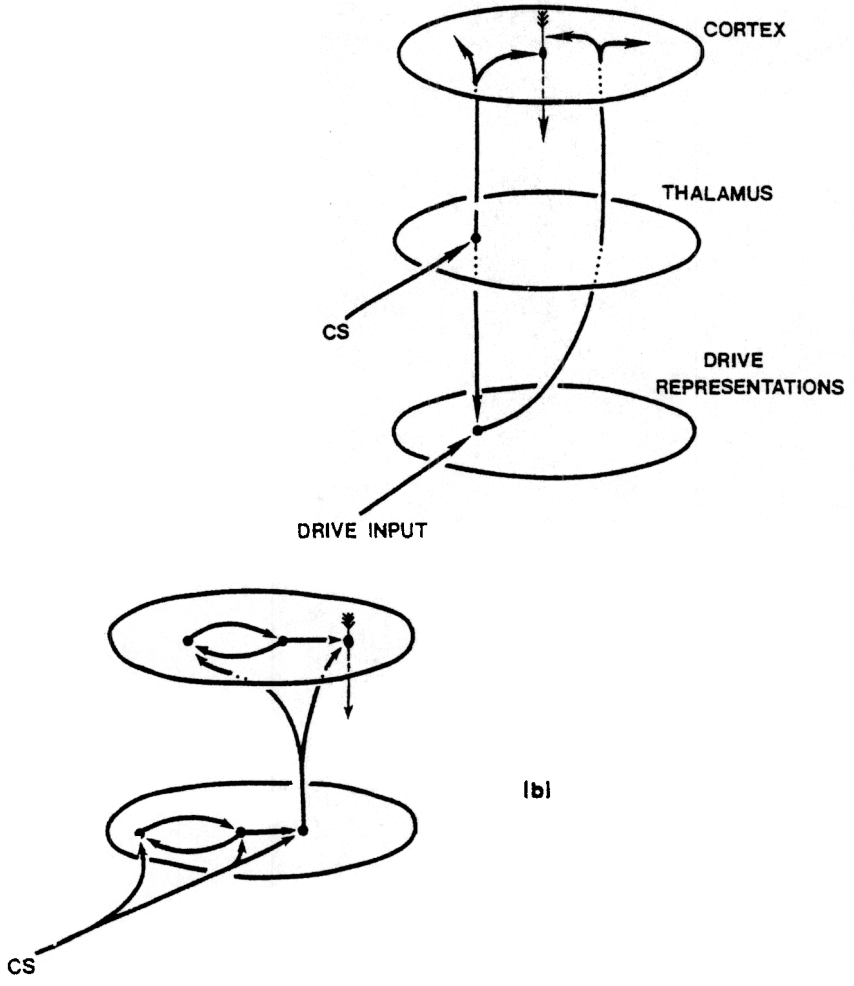

Fig. 15. (*a*) CS input excites nonpolyvalent thalamic cells which send signals both to neocortex and drive representations; (*b*) Intrathalamic and/or cortical reverberatory interneurons are now distinct from the nonpolyvalent output cells $v_{i1}^{(2)}$.

These anatomical constraints can help students of thalamocortical dynamics to correlate a given species variation with the corresponding network variation without losing sight of the fact that the entire class of networks can compute qualitatively similar functional transformations. An experimental framework wherein these anatomical alternatives may be testable has been developed by Gabriel *et al.* (1980), who have studied stimulus-reinforcement contingencies that are controlled by hippocampal interactions with cingulate cortex and anteroventral nucleus of the thalamus.

25. Motivational set: is incentive motivation conditionable?

Postulates (C) and (D) contain the main content of the synchronization problem and impose the basic network pathways with which we have been working. Once this formal framework is before us, it takes on a life of its own in two senses. First, it becomes clear that the network, as it stands, does not have certain important properties, but it is also easy to modify the network so that it does have these properties. Second, the overall structure of the network possesses a discernible symmetry. Improving this symmetry modifies the network so that it possesses new behavioral properties. The first and second procedures, moreover, both lead to the same modifications. Why are these modifications necessary? The answer is that there are other organizational principles than the synchronization problem at work *in vivo*. The synchronization problem fortunately implies enough network structure to force us into simple examples of these other principles, and in fact this was one route whereby the principles were discovered. This route is followed here because it is the most efficient way to derive mechanisms of reinforcement and attention. In Grossberg (1980*a*), these properties are derived from principles concerning the development of cognitive codes, which imply synchronization properties as a special case.

In Section 13, the symmetry properties of the network suggested that drive inputs exist. Postulate (F) gave this formal observation behavioral meaning by recognizing the need for a satiety mechanism. Now the symmetry properties suggest another network addition. The nonspecific conditioned reinforcer pathways from external cue representations to internal drive representations are conditionable. By contrast, the nonspecific incentive motivational pathways from internal drive representations to external cue representations are not. Should they be? Given the interpretation in Section 23 of the incentive motivational pathways as a CNV substrate, this question becomes: Can the CNV be directly conditioned? Are the apical dendrites of neocortical pyramidal cells the locus of this conditioning process?

If incentive motivation is conditionable, then we have at our disposal a mechanism for establishing a subliminal motivational set. After conditioning occurs, a CS could excite its drive representation via conditioned reinforcer signals. Then the internal drive representation could deliver incentive motivational signals preferentially to those external cue representations with which it was previously associated. In this way, activating a given external cue representation could sensitize an ensemble of motivationally related external cue representations via incentive motivational feedback. The sensitized representations form the subliminal motivational set.

The conditionability of incentive motivation is a necessary condition for avoiding some unpleasant behavioral properties. All of these properties are a result of the fact that incentive motivation is nonspecific. How can bad properties arise from this fact? We seemed to require that arousal be nonspecific to solve the synchronization problem in the first place (Section 5). I am not denying this basic

insight. It implies, and I reaffirm, that nonspecific arousal inputs are all initially strong enough to fire polyvalent sensory cells when they converge with CS inputs. In other words, the LTM traces across the nonspecific arousal pathways are all strong enough to cause large gated signals in response to arousal inputs. Still otherwise expressed, the pattern of LTM traces across the nonspecific arousal pathways is initially quite uniform and each LTM pathway is viable. I now suggest that conditioning can change this uniform LTM pattern by differentially strengthening the LTM traces that abut those polyvalent representations at which CS signals and incentive signals simultaneously converge during learning trials.

The unpleasant properties include the following examples. Suppose that two conditioned reinforcers, CS_1 and CS_2, are turned on simultaneously. Let each reinforcer preferentially project to a different drive representation D_1 and D_2, respectively. Let the drive input to D_1 be zero but the drive input to D_2 be large. Also suppose that CS_2 has not been conditioned to a motor habit. Cue CS_1, by itself, could only activate the first stage of its sensory representation, since the drive input to D_1 is too small to release incentive motivation from D_1. Cue CS_2, by itself, could only elicit an internal emotional reaction compatible with its drive representation. By contrast, if both cues are presented and the incentive pathways are not conditionable, then CS_2 can motivate performance of CS_1's habit, because D_2 will deliver incentive signals nonspecifically to all polyvalent cells, including those which represent CS_1.

In another version of the same dilemma, CS_1 and CS_2 are both conditioned to motor habits, the drive input to D_1 is again zero whereas the drive input to D_2 is large, but the intensity of CS_1 is large whereas that of CS_2 is small. Since CS_2 delivers equal incentives to the polyvalent cells of CS_1 and CS_2, the habit corresponding to CS_1 is favored because of CS_1's larger intensity despite the occurrence of zero drive input to D_1.

Another unpleasant consequence of unconditionable incentive is this. When we consider language behavior, we will want to understand how an internal need, such as hunger, can initiate an external language communication like 'I am hungry. What is there to eat?' This cannot be done in the present framework if the drive representations project uniformly to all cortical representations.

Consequently, I make the following postulate:

G. A given incentive can be associated with any of several external cue representations.

The minimal realization of postulate (G) is to suppose that incentive motivational signals are conditionable, and thus that subliminal motivational sets can be learned. This conclusion can be summarized in a fancier language that takes on important meaning when one studies the development of cognitive codes: conditioned reinforcer pathways and conditioned incentive motivational pathways are dual pathways in a network feedback module.

26. Distinct cortical recurrent loops for STM and LTM

Each formal constraint on network design needs to be realized in a physically plausible way. In the present instance, we need to ask: How can incentives be conditioned in the cortical analogs of Figure 14? How can the arousal pathway to the apical dendrites of pyramidal cells be conditioned at times when the pyramidal cells fire? In particular, how do the synaptic contacts at the apical dendrites know when the pyramidal cell body is firing?

There exist two possible answers. One suggests that an intracellular conditioning pathway exists, and the other that an intercellular conditioning pathway exists. The intracellular answer posits that cell body firing activates antidromic action potentials, or other signals, that invade the dendritic tree and condition whatever apical dendrite synapses are active at the time (Figure 16*a*). The intercellular

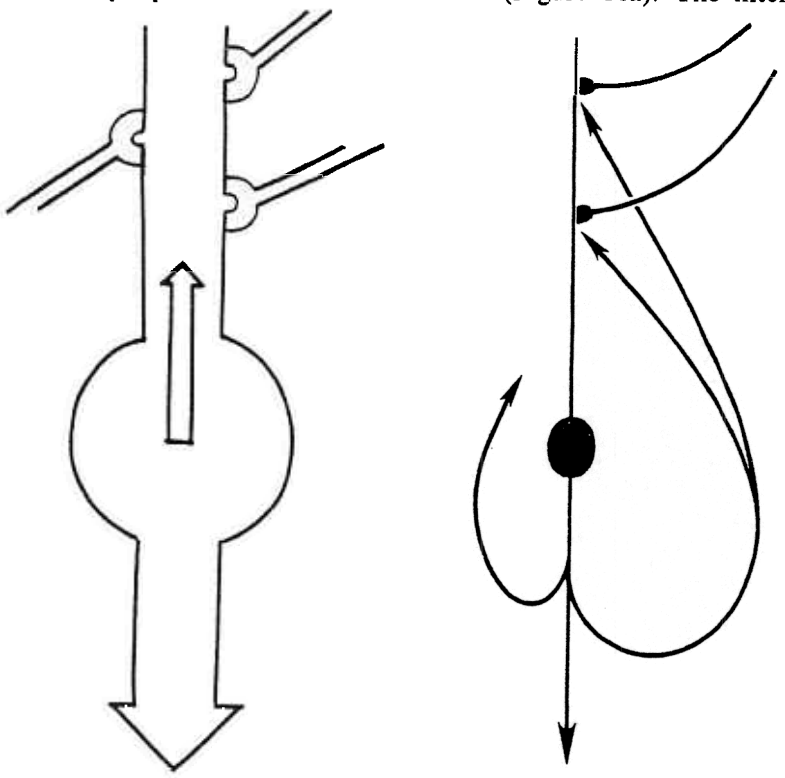

Fig. 16. (*a*) Cell body spiking might trigger massive antidromic action potentials that drive post-synaptic conditioning effects of simultaneously active synapses; (*b*) Cell body spiking might activate feedback pathways which drive conditioning of simultaneously active synapses by a presynaptic gating action.

answer suggests that active pyramidal cell axons also excite recurrent axon collaterals. These collaterals activate interneurons that terminate at the arousal cell–apical dendrite synapses and cause a conditioned change via a shunting mechanism at those apical dendrite synapses which are active (Figure 16*b*).

Both of these mechanisms suggest that cortical conditioning is driven by the suprathreshold activity of the pyramidal cells. Subthreshold activity is insufficient. Moreover both mechanisms include the pyramidal cell in excitatory recurrent interactions that are prerequisites for LTM storage, and that are distinct from the excitatory recurrent interactions that subserve STM storage.

These observations lead to several experimental questions, Can the CNV be conditioned if the pyramidal cells are prevented from firing? Under what circumstances *in vivo* do antidromic action potentials invade apical dendrites? Do synapses of cortical interneurons terminate on the synaptic knobs or dendritic spines of extracortical afferents to the apical dendrites? Do these interneurons get activated by pyramidal cell axon collaterals?

27. Motivation-dependent responses to sensory cues: multiple sensory representations or developmental competition for synaptic sites?

This section introduces a refinement of network design that satisfies another behavioral postulate in a minimal way.

In the network as it stands, once a sensory representation gains control over an incentive pathway and a habit, it cannot be used to learn another habit which is motivated by a different incentive pathway. For example, a visual cue at a choice point could elicit a left turn to get food, but it could not elicit a left turn to get food when hungry and a right turn to get food when thirsty. Of course, cues to the left or the right of the choice point could be preferentially associated with one or the other drive representation. Here we consider how a single cue could be differentially influenced by more than one drive. The issue of what cue combinations have internal representations will not be considered here. In Grossberg (1978*b*, Sections 25–47) I discuss how context-dependent internal representations, notably representations which are sensitive to particular sequences of events, can be generated. Herein I ask how any such representation can be differentially influenced by more than one drive. This property may not exist in all species.

The property in question can be achieved by allowing every sensory cue to be represented in several subregions each of which receives incentive pathways preferentially from a different drive representation. One way to do this is to let the sensory cues excite multiple sensory representations (Woolsey, 1958; Zeki, 1974) that are laid out in distinct network (e.g. cortical) regions. Another way is to suppose that, at the time when incentive pathways from the drive representations are developing, they compete with each other for synaptic space at each sensory representation, much as ocular dominance columns develop in the visual system (Hubel and Wiesel, 1977), or as corticostriatal terminals become fenestrated (Goldman-Rakic, 1981). As a result of this competition, the cells in each sensory

representation will be parcelled out into cell groups which receive more incentive pathways from one drive representation than any other. Here multiple incentive sensory representations do not exist, but within each sensory representation, the incentive pathways from different drive representations are clustered into distinguishable bundles. Looking over the entire cellular tissue, one would discern a patchwork quilt of overlapping sensory and drive sensitive areas. This configuration, should it exist, might be an evolutionary precursor of networks in which multiple sensory representations have been fully elaborated by a combination of synaptic competition abetted by a cell sorting process that segregates cells that become committed to particular drive representations into distinct sensory representations (Steinberg, 1970).

The behavioral constraint that yields this network refinement, where it exists, is summarized by the following postulate.

H. A discriminative cue can elicit distinct responses in different motivational contexts.

Multiple visual and auditory sensory representations are known to exist in vertebrate cerebral cortex (Woolsey, 1958; Zeki, 1974) but their relationship to drive-dependent response elaboration seems not to have been investigated. An effort should be made to study the patterns of axonal degeneration of drive representations in sensory areas and to correlate these patterns with that species' ability to discriminate cues in different motivational contexts. For example, what is the degeneration pattern of a drive region that supports enhanced self-stimulation when the animal is hungry? Sexually motivated?

28. Sensory-drive heterarchy: competitive decisions after drives and conditioned reinforcers interact

Typically, several external cues and drive representations are simultaneously active. How does the network decide which cues will be capable of eliciting observable behavior? What are the rules for parallel processing of cues? These general questions can be refined into a series of specialized issues which eventually force us to study new principles of network design. One of these principles leads to mechanisms whereby populations compete with each other (Grossberg, 1970*b*, 1973, 1980*b*). Sometimes the competition is organized in a feedforward anatomy (Figure 17*a*), but if the results of the competition are also stored in STM, a feedback anatomy is used (Figure 17*b*), whose feedback loops can store the pattern after its generating inputs terminate.

What keeps the network from simultaneously releasing two or more motivationally incompatible behaviors, such as eating and copulation? This question, phrased affirmatively, becomes our next postulate.

I. Motivationally incompatible behaviors compete.

Some form of competition between network channels is needed to sense the momentary balance of drives and available cues, and to decide which behavior is

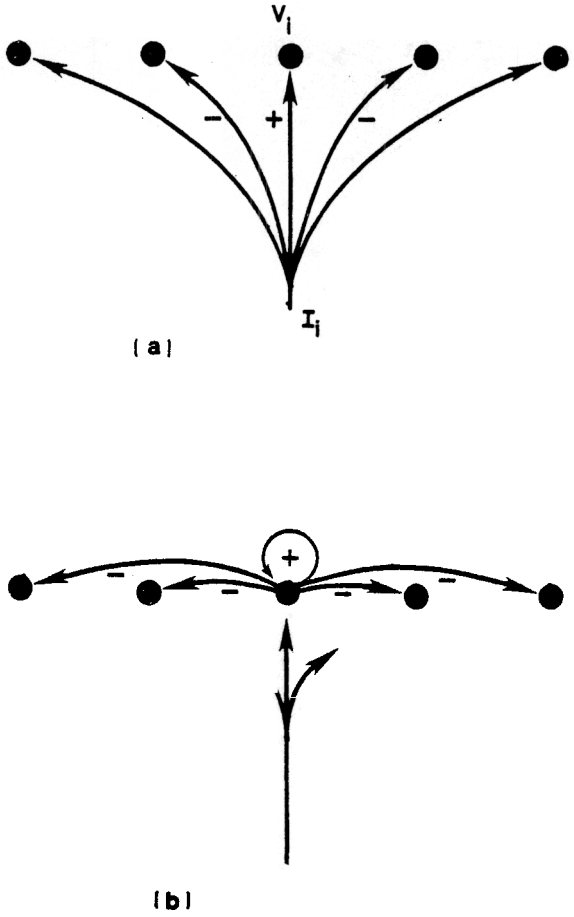

Fig. 17. In (*a*), a feedforward on-center off-surround network allows inputs to compete before generating outputs. In (*b*), a feedback competition allows the winning populations to have their activities stored in STM.

most appropriate at any time. Below a simple example of feedforward competition illustrates that this can be done. I will also indicate why feedforward competition is insufficient. It will then be shown that the anatomical stage at which competition acts must be carefully programmed in the network. If the competition occurs one synapse too soon, the network could not possibly survive.

To fix ideas, consider the following example of competition. Let n cell populations $v_1, v_2, \ldots, v_n$ be given and suppose that each population v_i is excited by a fluctuating input $I_i(t)$. Suppose that input $I_i(t)$ also inhibits all the populations v_k, $k \neq i$, with the same strength that it excites v_i. Let all the inputs, both excitatory and inhibitory, summate at each v_i, and let the activity $x_i(t)$ of v_i decay back to equilibrium at rate A when no inputs occur. Then

$$\frac{d}{dt}x_i = -Ax_i + J_i \tag{8}$$

where J_i is the net input

$$J_i = I_i \quad \sum_{k \neq i} I_k \tag{9}$$

to v_i. To see how the competition works, define the total input $I = \sum_{k=1}^{n} I_k$ and the relative input sizes $\theta_i = I_i I^{-1}$. Since each $I_i(t)$ can fluctuate wildly through time, so too can $I(t)$ and each $\theta_i(t)$. However, each J_i can be written as

$$J_i = 2I_i - I \tag{10}$$

which is the same as

$$J_i = 2I(\theta_i - \tfrac{1}{2}). \tag{11}$$

Since the θ_i's sum up to 1, at most one θ_i can exceed $\frac{1}{2}$ at any time. By (11), at most one J_i can be positive at any time. By (8), the v_i corresponding to this J_i is excited, whereas all other v_j are inhibited. Thus no matter how wildly the inputs fluctuate, the competition uses a majority rule to choose a definite winner at any time. The mechanism in (8) and (9) is called an additive feedforward on-center (excite v_i) off-surround (inhibit all v_k, $k \neq i$) network.

Not all networks of this type will perform a simple majority rule computation. For example, in the network

$$\frac{d}{dt}x_i = -Ax_i + I_i - \sum_{k=1}^{n} I_k B_{ki}, \tag{12}$$

the net inhibition of v_i by I_k is $I_k B_{ki}$, where B_{ki} measures the strength of the (k,i)th inhibitory pathway. The B_{ki} can easily be chosen so that more than one v_i can be excited at any time. Majority rule competition occurs if the inhibitory signals of the off-surround are broadly distributed actoss the field of oppulations.

These details are not our main concern now. The reader can, however, readily notice several basic deficiencies of majority rule competition, and more generally of feedforward competitive laws. For example, what happens if no θ_i exceeds $\frac{1}{2}$? Does the network simply do nothing? How can a network be designed that can retune its own sensitivity until a winning channel is found? What happens if the conditioned reinforcer inputs fluctuate very quickly? Does the motivational baseline also fluctuate because each J_i does? How can fluctuating cue and drive inputs be translated into a steady motivational baseline during the performance of each motivated act? What if inputs fluctuate so fast that J_i switches from positive to negative before its act can be performed? How can enough inertia be built into the competition to permit an act to be performed before motivation switches to another incentive channel? These are the types of crucial questions that can only be answered by mathematical analysis, because they all depend on surprising properties of competition in feedback networks (Grossberg, 1973, 1975, 1981*b*, 1982*a*,*b*,*c*).

For now, we content ourselves with considering the following question. At what stage of network processing does the competition act? There are two main alternatives, and one of them leads to disaster.

The first alternative allows the drives to compete among themselves before they are acted upon by conditioned reinforcers (Figure 18*a*). The second alternative lets drive inputs interact with conditioned reinforcer inputs at the drive repre-

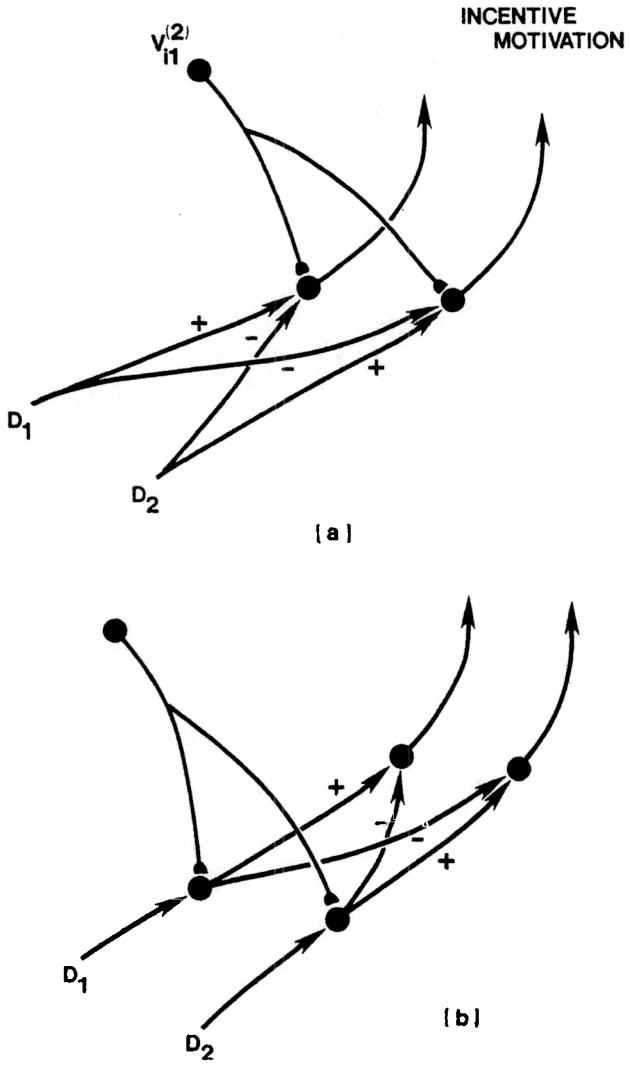

Fig. 18. In (*a*), a drive hierarchy prevents incentive motivation from being released except by the prepotent drive. In (*b*), a nonprepotent drive input can elicit incentive motivation if it is augmented by sufficiently strong conditioned reinforcer inputs.

sentations before the drive representations compete among themselves (Figure 18*b*). The first alternative is called a drive hierarchy and the second alternative is called a sensory-drive heterarchy for the following reason.

In a drive hierarchy, only the largest, or prepotent, drive can deliver a positive input to its drive representation. If no external cue is present that is compatible with this drive, then no incentive motivation can be released. The disaster is this. Even if external cues are available that are compatible with large, but nonprepotent, drive inputs, none of these cues can trigger observable behavior.

The sensory-drive heterarchy overcomes this dilemma. Here any drive representation can fire if it receives conditioned reinforcer and drive inputs. Then the active drive representations compete among themselves to decide which one will release incentive motivation. If no external cues compatible with the prepotent drive are available, then its drive representation does not fire, so it does not prevent a drive representation with a smaller drive input from winning the competition. Moreover, if two drive representations $\mathscr{D}_1$ and $\mathscr{D}_2$ both receive positive drive inputs D_1 and D_2 respectively, such that $D_1 > D_2$, nonetheless $\mathscr{D}_2$ can win the competition if its conditioned reinforcers are more active than those of $\mathscr{D}_1$. In all, the sensory-drive heterarchy computes which combination of available cues and drives is dominant at any moment, not just which drive is dominant at any moment.

To illustrate the heterarchical concept, I will describe some rules that simplify real-time computations which occur in competitive feedback networks. To start, let us suppose that drive inputs D_i and conditioned reinforcer inputs S_i combine multiplicatively at $\mathscr{D}_i$. This rule can be interpreted in several ways. It formally says that both factors are needed to fire $\mathscr{D}_i$. It can be physiologically interpreted by saying that drive inputs sensitize drive representations to conditioned reinforcer inputs, or that drive inputs gate conditioned reinforcer inputs. Although we formally write products like S_iD_i below, in a finer physiological description, cues and drives do not play a symmetric role. A simple version of the heterarchical idea is the rule: $\mathscr{D}_i$ fires only if

$$S_iD_i > \max(\varepsilon, S_kD_k : k \neq i) \tag{13}$$

where ε is a threshold that must be attained before any drive representation can fire. By (13), only one drive representation can fire at any time; namely, that $\mathscr{D}_i$ whose gated cue inputs S_iD_i are maximal. This law builds some temporal stability into the delivery of incentive motivation, since inequality (13) can persevere throughout a time interval during which the individual cue signals S_j are fluctuating wildly. The simplest law whereby a stable baseline of motivation is also achieved follows readily from (13). Let $x_i(t)$ be the activity of $\mathscr{D}_i$ at time t. Suppose that all the drive representations remain inactive, or at best subliminally active, until (13) holds at some $\mathscr{D}_i$. Then x_i is rapidly amplified by feedback competitive interactions until it attains the value $B > 0$, which is the suprathreshold motivational baseline. All other x_j are quickly driven to zero by the recurrent competition, and these motivational values are then stored in STM while (13) holds. In mathematical terms,

$$x_i = \begin{cases} B, \text{ if } S_iD_i > \max(\varepsilon, S_kD_k: k \neq i), \\ 0, \text{ otherwise.} \end{cases} \tag{14}$$

The rule (14) is the simplest law that builds some temporal stability and a motivational baseline into the heterarchical computation. This rule also overcomes a disadvantage of majority rule competition: Just so long as some S_iD_i exceeds the threshold ε, which can be a small parameter or even zero, a motivational choice is made, and one drive representation achieves the baseline value B. Figure 19 schematizes an anatomy that can compute this competitive rule.

The performance of a competitive feedback network is more complex than equation (14). For example, such networks can maintain any of a finite, or infinite, set of operating levels even after their inputs are shut off (Grossberg, 1973, 1978*c*, 1980*a*,*b*). The inputs can determine, by their own size, which operating level will be chosen. Moreover, inputs that are left on can modify the network's possible

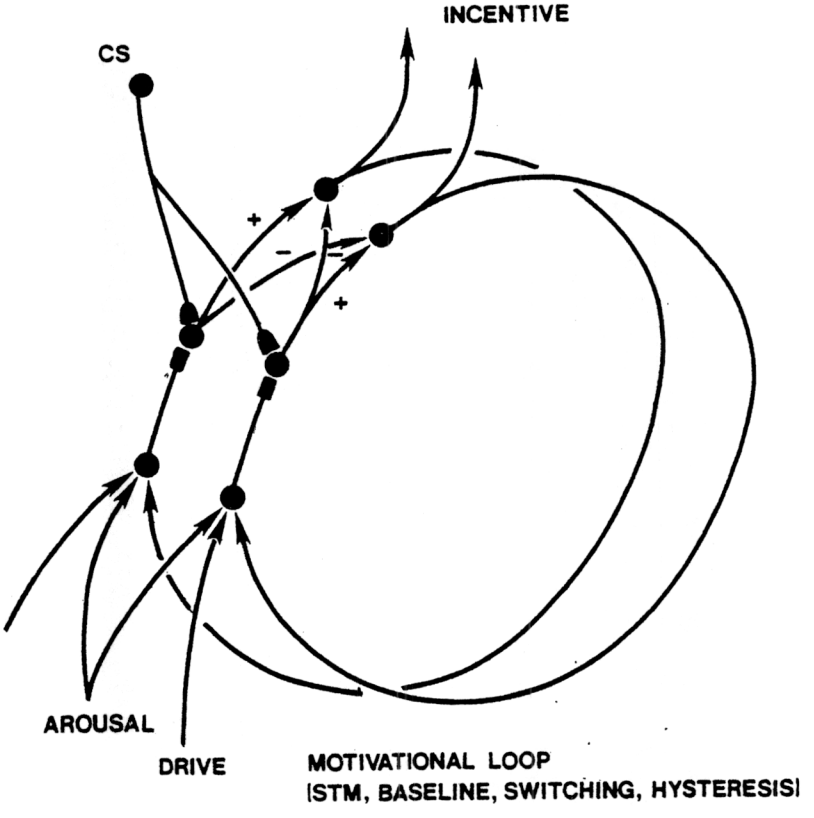

Fig. 19. A feedback competitive network whose signals are modulated by slow transmitter gates can maintain a steady motivational baseline, and can switch between motivational channels in response to sufficiently large changes in conditioned reinforcer, drive, or arousal inputs.

operating levels. These subtleties will not be needed to draw our main conclusions.

Even the rule (14) is deceptively simple, because S_i is the net effect of all conditioned reinforcers on $\mathscr{D}_i$. If $I_m b_{mi}$ is the output from the mth sensory representation $v_{mi}^{(2)}$ to $\mathscr{D}_i$, and z_{mi} is the LTM trace in the pathway from $v_{m1}^{(2)}$ to $\mathscr{D}_i$, then

$$S_i = \sum_{m=1}^{n} I_m b_{mi} z_{mi}. \tag{15}$$

Substituting (15) into (14) yields

$$x_i = \begin{cases} B, \text{ if } \sum_{m=1}^{n} I_m D_i b_{mi} z_{mi} > \max(\varepsilon, \sum_{m=1}^{n} I_m D_k b_{mk} z_{mk} : k \neq i), \\ 0, \text{ otherwise.} \end{cases}$$

Moreover, z_{mi} summarizes the entire reinforcement history between the mth internal representation $v_{m1}^{(2)}$ and $\mathscr{D}_i$ by correlating all past occurrences of I_m and x_i. The simplest such law (Grossberg, 1964, 1968, 1971*b*, 1972*a*) says that z_{mi} computes a time-average, with constant decay rate c, of the product of signal I_m and postsynaptic potential x_i. Then

$$\frac{d}{dt} z_{mi} = -cz_{mi} + d_{mi} I_m x_i$$

which is the same as

$$z_{mi}(t) = z_{mi}(0)e^{-ct} + d_{mi} \int_0^t e^{-c(t-v)} I_m(v) x_i(v) dv. \tag{18}$$

Equations (16) and (18) together summarize the simplest heterarchical computation. These equations illustrate how an animal's present decisions, as in (16), can reflect all of its past reinforcement, drive, and motivational history, as in (18). In particular, z_{mi} in (18) correlates I_m with x_i, where I_m connotes the momentary strength of a CS or UCS input, and x_i connotes the momentary strength of a motivational variable that is determined by all the present drives and cues, as well as all the prior drives and cues that ever influenced the network. Thus the conditioned reinforcer LTM traces feed upon themselves: They guide the present motivational decision, which is thereupon sampled and alters their LTM values. Perhaps the most important thing to keep in mind about (16) and (18) is not all the subtleties of these feedback effects, but rather that the formalism summarizes so much subtlety in just two equations.

The heterarchical concept elegantly explains how an animal's observable behavior can seem to be unrelated to a specific drive until an appropriate releasing cue is presented (Bolles, 1967): A high drive D_i can occur while $x_i = 0$ if another drive representation $\mathscr{D}_j$ wins the heterarchical competition. The heterarchical concept also clarifies some differential effects of parametric changes in drive and reinforcement on learning rate versus performance speed in the following way.

29. Differential effects of drive and reinforcement on learning rate versus performance speed

An increase in drive level can affect the probability that a rat will run down an alley to be rewarded, its reaction time to run, and the vigor of the running response (Estes, 1958; Mackintosh, 1974). In particular, increasing drive decreases the likelihood that the rat will indulge in competing behaviors (Cotton, 1953). This fact can induce a complex interaction between drives and learning, since an animal who runs only a little before competing responses interfere can learn a different response series than an animal who continues running down the alley (Bolles, 1967; Campbell and Kraeling, 1953). If these momentary effects of drive on performance are synthesized into a behavioral plan, then different running responses can be engendered on later trials, even if the drive level is changed.

To see how drive can alter probability and reaction time of running, consider (16). In (16), an increase of D_i increases the probability that, and decreases the reaction time with which $\mathscr{D}_i$ will win the heterarchical competition. Once $\mathscr{D}_i$ wins, $x_i = B$ no matter how large D_i is. In other words, increasing D_i can reduce the reaction time and increase the probability of running, but not necessarily alter the motivational baseline that supports the running response. To understand how the motivational level, and the induced running speed, can depend on drive input and conditioned reinforcer input size, one needs to study the operating levels B of competitive feedback networks in some detail. This study will not be given here. Instead, I will indicate how the simplest heterarchical rules help to explain various other data to lend further support to the heterarchical concept.

The following paradoxical finding can, for example, be explained. Animals that are trained on high drive can maintain their high drive performance characteristics when they are switched to low drive (Capaldi, 1971; Mackintosh, 1974; Zaretsky, 1966), but not conversely (Bolles, 1967; Desse and Carpenter, 1951). The perseverative effect of high drive performance during low drive suggests that high drive influences performance through a learned effect. Analogously, a large value of D_i in (16) on learning trials increases the probability that x_i wins the competition. By (18), those LTM traces z_{mi} with active cues I_m will attain large values. These cues thereupon enjoy powerful conditioned reinforcer properties which can motivate the animal to run. On low drive trials, these cues are still available, their inputs I_m are still large, and they are still gated by large z_{mi} values that were learned under high drive. Consequently the gated inputs $I_m z_{mi}$ remain large after the switch to low drive trials. Suppose that D_i is decreased, but is large enough for $\mathscr{D}_i$, with the help of large terms $I_m z_{mi}$, to still reliably win the heteratchical competition. Then the high drive performance level will persist into the low drive situation because it is bolstered by conditioned reinforcer inputs. If D_i can regularly win the competition in this way, then the LTM traces z_{mi} continue to be boosted on low drive trials if they follow high drive trials. The converse transition, from low drive to high drive, does not show persistence. For example, given low drive levels at the outset, $\mathscr{D}_i$ cannot reliably win the competition, so little learning occurs. An

increase of drive then rapidly causes a higher performance level to occur because it directly changes the competitive balance.

Amount of reinforcement can influence rate of learning as well as asymptotic learning speed (Mackintosh, 1974). Moreover, a change in reinforcement level from high to low, or conversely, can quickly change the animal's response rate (Crespi, 1942; Mackintosh, 1974; Zeaman, 1949) by contrast with the effects of changing drive level.

Why do parametric changes in drive and reinforcement level have different effects on performance? Several factors can work together to produce these properties. One such factor is the following. A large reinforcement on each trial creates a large cue input I_m. The cue input directly excites the polyvalent sensory cells, unlike drive inputs, as well as indirectly exciting the polyvalent sensory cells via incentive feedback to these cells. The larger reinforcer's cue properties hereby increase the readout of its habit on a moment-by-moment basis. Consequently, a sudden change in reinforcer level can cause rapid performance changes, even if the motivational feedback itself has only one suprathreshold level. This explanation of reinforcer effects depends on the existence of variable I_m values. Changes in reinforcement that do not change I_m do not have these properties. For example, an increased concentration of sucrose or saccharin in a liquid reward reliably has these effects.

Reinforcer effects can also be influenced by the action of disconfirmed expectancies. To illustrate this possibility, I will use properties that are discussed in Sections 32–34, below. These remarks can be skipped on a first reading.

An animal who expects a high probability of reinforcement can experience a negative incentive motivational rebound on a trial during which an expected reward does not occur because the reinforcement probability has been reduced (Grossberg, 1972*c*, 1975). The rebound has the same effect as an increase in negative drive input even though no change in internal drive has occurred. The rebound can hereby rapidly reduce the net motivational support for the behavior and reduce the animal's response rate. In a similar fashion, when a low probability of reinforcement is expected, a positive incentive motivational rebound can supplement direct consummatory inputs on a trial when the animal is unexpectedly rewarded due to an increase in reinforcement probability. A rapid increase in motivational support is hereby achieved. Transient overshoots in incentive motivational outputs can hereby occur due to unexpected changes in reinforcing events. These transients can reflect themselves in transient behavioral contrast effects (Mackintosh, 1974). Sustained behavioral contrast effects can be produced after the transient effects are encoded in LTM by the conditioned reinforcer and incentive motivational pathways. This analysis differs from the one that Hull (1952) and Spence (1956) gave to explain these effects. These authors assumed that reinforcement decrement directly influences incentive level. I claim that the direct effects are specific cue effects, including a lower probability of winning the heterarchical competition if the reinforcement decrement is sufficiently large. Indirect effects include rebound effects that are triggered by rapid changes in cue level or by unexpected events.

30. Drives and reinforcers: multiplicative, additive, or neither?

Hull (1952) claimed that reinforcement and drive interact multiplicatively, whereas Spence (1956) claimed that they are independent, or interact additively. Experiments to test this distinction have tended to yield mixed results. Some experiments suggest additivity (Reynolds and Pavlik, 1960; Pavlik and Reynolds, 1963). If either the drive level or the reinforcement level is assigned a very low value, however, a significant multiplicative interaction is found (Seward *et al.*, 1958; Seward *et al.*, 1960).

The rule (16) sheds some light on this controversy. In (16), drives D_i and cues S_i do interact multiplicatively. However, the competitive law also contains additive elements. For example, given a fixed S_i level, the competition computes whether $D_i \geqslant \varepsilon S_i^{-1}$, which is an additive effect. Similarly given a fixed D_i level, the competition computes whether $S_i \geqslant \varepsilon D_i^{-1}$, which is again an additive effect. Moreover, given fixed levels S_i and S_j, the competition computes the sign of $S_iD_i - S_jD_j$ which is a linear combination of D_i and D_j, again an additive effect. The same is true if D_i and D_j are fixed while S_i and S_j vary. Both additive and multiplicative elements enter the heterarchical computation. The multiplicative effects are rate-limiting at very small values of S_i and/or D_i, since no matter how large S_i is, if $D_i \simeq 0$ it follows that $S_iD_i \simeq 0 < \varepsilon$, and no matter how large D_i is, if $S_i \simeq 0$ it follows that $S_iD_i \simeq 0 < \varepsilon$.

A final comment on the motivational baseline is in order. The existence of such a baseline illustrates how a recurrent network can store an activity level in STM despite the existence of sufficiently small temporal perturbations in its inputs. This is not, however, all the recurrent networks can do. Their overall operating levels can be influenced by the size of specific inputs as well as by nonspecific shunting interactions, so that *in vivo* the motivational level need not be constant through time. Our main point has been that these operating levels can defend themselves against momentary fluctuations in cues or drives until a hysteresis boundary is reached, after which a switch to a new level can be rapidly effected.

31. Suppression by punishment

The discussion has thus far focused on how an animal can learn consummatory behavior. It has not considered how an animal can rapidly modify or terminate consummatory behavior when changing environmental contingencies render the behavior maladaptive. In particular, we have not considered how an aversive cue can suppress behavior. Nor have we considered how the nonoccurrence of an expected consummatory goal object can suppress behavior. Now the minimal mechanisms of aversive conditioning, whether by punishment or frustration, will be reviewed. I will conclude that the mechanism whereby a punishing event suppresses behavior, via classical conditioning, automatically possesses properties which enable the unexpected nonoccurrence of a desired goal object to extinguish behavior, via instrumental conditioning. This important fact, which sets my theory

apart from other efforts to model reinforced behavior, follows from the gating properties of chemical transmitters in competitive anatomies (Grossberg, 1972*c*, 1981*a*, 1982*a*,*b*,*c*).

Many experiments about punishment and avoidance are thoroughly reviewed elsewhere (Bolles, 1969; Dunham, 1971; Estes, 1969; Grossberg, 1972*b*, 1975; Mackintosh, 1974; Pearce and Hall, 1980). A basic theme running through this literature elaborates the fact that punishment can suppress goal-oriented behavior without extinguishing knowledge of the goal. Here, I will only be interested in showing how a punishment and avoidance mechanism easily flows from the ideas which are already at our disposal. Various subtle conditioning properties—such as self-punitive behavior, overshadowing, superconditioning, learned helplessness, peak shift and behavioral contrast, partial reinforcement acquisition effect and novelty as a reinforcer—have elsewhere been derived as properties of the mechanisms in prescribed input environments (Grossberg, 1972*b*,*c*, 1975, 1981*a*, 1982*a*,*b*).

Our point of departure is the fact that CS-activated sampling cells are polyvalent (Section 5). This fact is a double-edged sword. It prevents the sampling cells from firing unless they simultaneously receive specific cue input and nonspecific arousal inputs. But it also forces the sampling cells to fire when these input conditions are achieved. Without additional network structure, these polyvalent cells must fire at disastrously unappropriate times as environmental contingencies change.

In particular, suppose that a cue CS_1 has been conditioned to a positive incentive motivational source and to a prescribed motor habit. Then whenever the drive subserving the motivational source is high, presentation of the cue will elicit the behavior because the polyvalent cells in its representation receive convergent specific and nonspecific inputs. This behavior will continue to occur no matter how unappropriate its consequences have become; e.g. if pressing a lever in response to the cue lever now elicits shock instead of food.

Some mechanism is needed to stop behavior that prior contingencies have started. In a theoretical vacuum, this factual triviality is not very constraining on possible theories. It becomes highly constraining, however, when it is recast as a formal design within the networks that have already been derived. We therefore ask: How can the polyvalent cells be prevented from firing? Can this be done such that the mechanisms for starting versus stopping behavior are symmetric?

A dual side of this issue exists. It is the following issue: Suppose that a mechanism for stopping unappropriate behavior can be found. What prevents this mechanism from stopping all future behavior? How can more appropriate behavior be selected despite the suppressive effects of this mechanism? For example, while an aversive cue is on, such as an intense shock, an animal might emit many erroneous escape behaviors before a random act succeeds in terminating the aversive cue. How can the correct escape act be learned despite the fact that the prior errors are not learned?

This observation can be rephrased as a postulate.

J^-. Onset of an aversive event can suppress new reinforced learning.
J^+. Offset of an aversive event can trigger new reinforced learning.

We ask for the minimal change in network design that can realize postulate (J^-). First we note that a change of LTM habit strengths in the pathways from sensory to motor representations cannot suffice for the following reasons:

(1) *Passive extinction.* If cell firing is prevented in these pathways for a long time, then the LTM traces might slowly decay. However, this process takes too long to prevent postulate (J^-) from being violated. Also there exist variations of the LTM law (17) in which no passive extinction occurs; for example, laws such as

$$\frac{d}{dt} z_{mi} = I_m d_{mi}(-z_{mi} + x_i) \tag{19}$$

wherein offset of the sampling signal $I_m d_{mi}$ from population v_m to v_i shuts off the rate of change dz_{mi}/dt of the LTM trace z_{mi}. Even where passive decay is possible, it can be retarded or reversed if recall trials intermittently occur while the animal is hungry. Then positive incentive motivational feedback can fire the sampling cells and thereby refresh, or restore, suprathreshold LTM levels via post-tetanic potentiation (Grossberg, 1969*c*).

(2) *Interfering habit.* Another possible way to disrupt the habit strength LTM traces is to suppose that the aversive event (e.g., shock) directly generates a motor command at the motor representations which interferes with the command read-out via the habit strengths. As the cue representations sample this new command, their habit strengths can encode it in LTM, rather than the previous consummatory command.

This mechanism also suffers from severe flaws. First, the network cannot learn specific avoidance tasks, since the shock stimulus rather than a specific avoidance response maintains the new motor command. Second, the network remains conditioned to the old drive representation, say a hunger representation. It can thus indulge in autonomic preparations for eating without being able to eat. Finally, the network remains maladaptively fearless, since only positive consummatory drives are conditional to external cues. The fearful meaning of the aversive event is nowhere represented in the network.

A mechanism is needed which can rapidly react to an aversive cue; viz., an STM mechanism. Since fast changes in motor commands do not suffice, we must consider fast changes due to the cue and drive properties of aversive events. These STM changes must be capable of driving slower LTM changes in conditioned reinforcer and incentive motivational pathways that can encode adaptive behavioral modifications.

The minimal drive property of an aversive event will now be introduced (Grossberg, 1972*b*,*c*). Let shock create an input at its own drive representation $\mathscr{D}_f$. Let this input be a monotone increasing function of shock intensity. Suppose that $\mathscr{D}_f$ can elicit signals which inhibit the positive incentive motivational outputs of positive drive representations. Also suppose that cue representations send

conditional pathways to $\mathscr{D}_f$ as well as to the other drive representations (Figure 20). Six important conclusions follow from this simple construction. (i) An intense

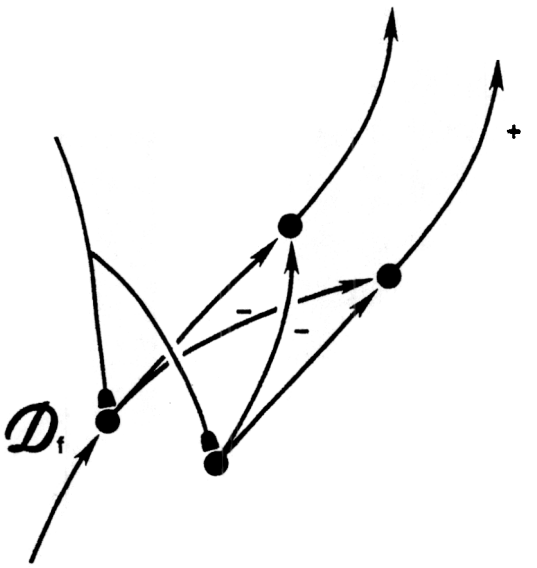

Fig. 20. When a drive representation whose activation is associated with the emotion of fear is included in the network and given negative incentive motivational properties, it helps to explain many properties of conditioned emotional responses.

shock can rapidly suppress consummatory behavior by inhibiting the motivating effects of positive reinforcers. The polyvalent sampling cells are hereby prevented from firing even when the first stages of their internal representations are firing. (ii) This suppression does not extinguish the LTM of the habits that are already encoded in the motor commands. It merely prevents these habits from being read out by polyvalent cell signals during unappropriate circumstances. (iii) This type of suppression can occur much faster than passive extinction of LTM traces. (iv) An intense shock can also prevent new habits from being learned by inhibiting release of sampling signals from polyvalent cells. (v) If an external cue is present only when the shock is on, its internal representation can learn a strong LTM connection to the $\mathscr{A}_f$ drive representation even though its polyvalent cells cannot fire. On later trials, onset of this cue can hereby elicit the suppressive (and emotional) effects of shock in the absence of shock. The cue hereby becomes a CS^+ that elicits a conditioned emotional response (Estes and Skinner, 1941). (vi) If a cue which has previously been conditioned to a positive motivational source remains on while shock is on, its LTM trace to the $\mathscr{A}_f$ drive representation can grow. Eventually, the cue's gated signals to the consummatory and $\mathscr{A}_f$ representations can become approximately equal, or the signal to $\mathscr{A}_f$ might even be larger. After the output from these drive representations compete, the cue's net motivational effect on polyvalent sampling cells approaches zero, or a negative value. The cue's

ability to motivate consummatory behavior has hereby been extinguished by competing signals from the $\mathscr{A}_f$ representation even if it continues to elicit a large conditioned output from the consummatory drive representation.

32. Antagonistic rebound and learned avoidance

To deal with postulate (J^+), we must again face the rigid law which governs polyvalent cell firing. When a shock input terminates, the input to $\mathscr{A}_f$ terminates. This event eliminates the direct source of suppression due to the shock. When this happens, the motor command (e.g., the lever press) which terminated the shock is active in STM, and the sensory feedback cues contingent upon this command (e.g., looking at the lever) have activated their internal representations. How can associative links between these active sensory and motor representations be learned? What motivational source enables the polyvalent cells of the sensory feedback cues to release sampling signals? What motivational source is sampled by the sensory feedback cues to endow them with conditioned reinforcer properties?

Our main dilemma is that offset of shock removes a source of negative input from the polyvalent cells' motivation. We need more, however, to fire a polyvalent cell. We need a source of positive motivation that is triggered by the offset of shock. This positive motivational source must be on only transiently after shock turns off. If it were on permanently, it could be used to motivate behaviorally irrelevant sensory-motor associations.

We therefore conclude that another arousal source exists. Speaking heuristically, this new arousal source supplies the motivational support for learning an avoidance response. Its activity is correlated with the internally perceived relief that can occur after offset of a fearful cue under appropriate experimental conditions (Denny, 1971; Masterson, 1970; Reynierse and Rizley, 1970).

Let us introduce a symmetric notation to highlight the relationship between the fear and relief representations. Denote by $\mathscr{D}_f{}^-$ the arousal cells that are excited by termination of shock input at the cells $\mathscr{D}_f$, which we henceforth denote by $\mathscr{D}_f{}^+$. The $\mathscr{D}_f{}^+$ cells are 'on' cells. They are turned on by shock, and remain on until shock is turned off. The cells $\mathscr{D}_f{}^-$ are 'off' cells. They are turned on temporarily by shock termination. On-cells and off-cells are familiar physiological components (Thompson, 1967, pp. 253, 349). Our theoretical task is to understand well enough how these components are designed to derive quantitative predictions about them.

The on-cells and off-cells are assumed to reciprocally inhibit each other to generate an incentive motivational output corresponding to fear, or to relief, but not both simultaneously.

The operation whereby offset of shock can elicit transient relief is called *antagonistic rebound.* The classical notion that instrumental reinforcement is due to drive reduction when shock terminates is replaced by rebound from negative-incentive motivational on-cells to positive-incentive motivational off-cells.

Similarly, offset of a positive-incentive motivational on-cell can generate a negative-incentive motivational off-rebound, as when a food source is suddenly withdrawn, even if no reduction of hunger drive has occurred.

At this point, we are very close to understanding the *gated dipole* mechanism on which all the more advanced properties of reinforcement, motivation, and expectational processes of the theory depend. This and the next section will sketch a few qualitative properties of the rebound concept which need no mathematics to be understood, and will motivate the derivation of the gated dipole theory.

The network must be expanded once again to allow cues to become conditioned to the new arousal source (Figure 21). Thus each sensory representation now sends

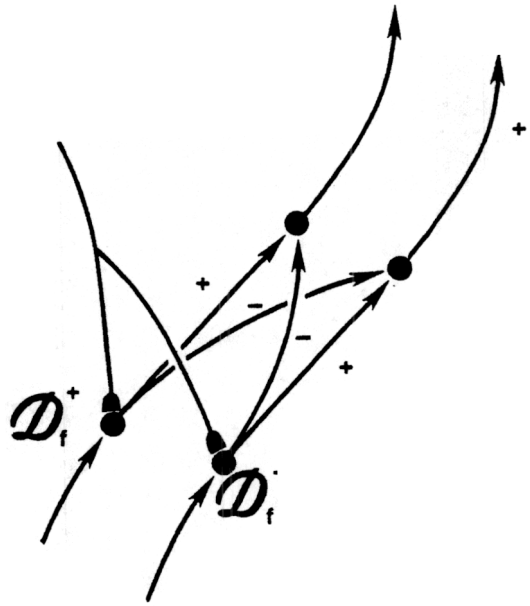

Fig. 21. A gated dipole in which fear and relief drive representations compete helps to explain various data about the balance between positive and negative reinforcement and extinction.

pathways (axons) to $\mathscr{D}_f^-$ as well as $\mathscr{D}_f^+$ and other arousal sources, such as the hunger representation $\mathscr{D}_h$. At any time, the synaptic knobs of each cue representation encode in their LTM traces a spatial pattern which represents a weighted average through time of all the motivational activities that the representation sampled while it was active. The net motivational output to polyvalent cells is determined by the inputs from all drive and cue representations. Even if half the cues send large conditioned signals to $\mathscr{D}_f^-$, no positive incentive will ge generated if the other half of the cues send equally large conditioned cues to $\mathscr{D}_f^+$. The competition between drive representations will annihilate these large inputs before

an output can be elicited. Similarly, shock termination yields little relief via antagonistic rebound if it is accompanied by onset of a conditioned aversive cue, as in a two-way shuttle box. Even shock termination is not necessarily rewarding in all environmental contexts. For the purpose of deriving the rebound mechanism, the section considers the simple case wherein a single aversive cue is turned on and off through time.

33. Slowly accumulating transmitter gates in tonically aroused competing channels cause antagonistic rebound

The heuristic argument leading to the dipole theory can be divided into eight steps.

(1) *Existence of a tonic input.* When shock terminates, the relief center $\mathscr{D}_f^-$ can emit a transient output. What input energizes this output? Offset of a shock input merely removes input activity. Some other input source must energize the relief rebound.

Since shock offset is the only change in external input, the input that energizes relief must be internally generated. Terminating shock somehow unmasks the effects of this internal input. The internal input to $\mathscr{D}_f^-$ is therefore not turned on or off by shock offset. It is also not turned off by shock onset, since then it would always be off. The internal input is therefore on throughout the learning experiment. It is a *tonic* input.

(2) *Existence of accumulation–depletion.* Output from the relief center is always transient. How is this accomplished? No externally driven input is available to accomplish this. The relief output is somehow depleted by its own activity. In other words, when shock is on, an accumulation process occurs in the relief center. After shock is turned off, the relief output is an increasing function of the amount accumulated at each time. This amount is gradually depleted until the relief output shuts off.

(3) *Competition between fear and relief.* We suppose that at most one of the outputs from $\mathscr{D}_f^+$ to $\mathscr{D}_f^-$ is nonzero at any time. In other words, either fear or relief, but not both, can be experienced by the network at a given time. Thus the final state of processing in $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$, before incentive motivational outputs are generated, is due to competition, or mutual inhibition, of the $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$ signals.

(4) *Existence of nonspecific accumulation–depletion in both channels.* When shock is off for a long time, outputs from both $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$ are zero. Thus the accumulation process at $\mathscr{D}_f^-$, driven by its tonic input, is balanced by a process going on at $\mathscr{D}_f^+$. The simplest idea is that a parallel process of accumulation–depletion, driven by its own tonic input that equals the $\mathscr{D}_f^-$ tonic input, takes place in the fear channel. In particular, the tonic input is delivered nonspecifically to both channels. It is therefore called a *nonspecific arousal* input. When shock is turned on, the shock input summates with the tonic input in the fear channel.

(5) *The rebound is slow.* It lasts at least seconds rather than milliseconds. It is a slow process compared to the fluctuation rates of cell potentials in response to input changes. After shock terminates, neither the fear nor the relief channel receives an externally driven input. Without an additional mechanism at work, their outputs would rapidly equalize, but they do not. Thus there exists a process slower than potential fluctuations that can bias output from $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$ in favor of $\mathscr{D}_f^-$ after shock terminates.

(6) *Both fear and relief are increasing functions of shock duration and intensity.* Both increasing duration (Annau and Kamin, 1961; Boe, 1966; Borozci *et al.*, 1964; Church *et al.*, 1967; Keehn, 1963; Strouthes, 1965) and intensity (Annau and Kamin, 1961; Boren *et al.*, 1959; D'Amato *et al.*, 1967; Huff *et al.*, 1967; Johnson and Church, 1965; Kamin *et al.*, 1963; Martin and Reiss, 1969; Reiss, 1970) can influence the strength of conditioned emotional responses and conditioned avoidance responses. These results suggest that both channels contain slowly varying processes which parametrically depend on shock intensity and duration when shock is on, and which counterbalance each other when shock is off for long intervals.

(7) *The relative balance of accumulation is changed by shock.* What causes the relief rebound to shut itself off? Is complete depletion of the accumulated product at $\mathscr{D}_f^-$ responsible for this property? Were this the case, then the tonic input alone could deplete $\mathscr{D}_f^-$, since it is the only input to $\mathscr{D}_f^-$. By symmetry, during shock, the shock input plus the tonic input could surely deplete $\mathscr{D}_f^-$. This does not occur, however, since the fear reaction is maintained by a long shock. A weaker conclusion is therefore necessary. Shock shifts the relative balance of accumulation in the two channels by depleting the fear channel more than the relief channel.

(8) *Signal size is a joint function of input size and amount accumulated.* This observation is crucial. During a relief rebound, both $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$ receive equal arousal inputs which ultimately balance the amounts accumulated at $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$, and thereby shut off incentive outputs. Before this can happen, $\mathscr{D}_f^-$ output exceeds $\mathscr{D}_f^+$ output because $\mathscr{D}_f^-$ accumulation exceeds $\mathscr{D}_f^+$ accumulation. In other words, given a fixed input size (the equal arousal inputs to $\mathscr{D}_f^+$ and $\mathscr{D}_f^-$), output is an *increasing* function of amount accumulated. This is true in each of the two channels.

When shock is on, increasing shock intensity increases $\mathscr{D}_f^+$ output, since it causes an increase in fear. Increasing shock intensity also *decreases* the amount accumulated at $\mathscr{D}_f^+$; this is the basis of rebound at $\mathscr{D}_f^-$ when shock is turned off. Thus output is not a function of accumulation level alone, since then increasing shock intensity would have decreased $\mathscr{D}_f^+$ output by decreasing the amount accumulated at $\mathscr{D}_f^+$. Output size is a joint function of input size (a fast variable) and accumulation level (a slow variable).

These arguments are sufficiently constraining to be capable of reducing our theory to rubble. When they were first made in Grossberg (1972*c*), however, it

was already apparent that the theory could quantitatively realize their demands with no difficulty. This was because the necessary laws had already been derived from associative learning postulates in Grossberg (1968, 1969*c*). For completeness, I review the simplest version of the transmitter gating rules that are needed to understand gated dipole dynamics. More complete derivations and analyses of these laws are found in Grossberg (1980*a*, 1981*a*, 1982*a*).

The simplest law whereby a chemical transmitter z gates an input signal S to yield an output signal T is

$$T = Sz, \tag{20}$$

where $z(t)$ obeys an accumulation–depletion law of the form

$$\frac{dz}{dt} = A(B - z) - T.$$

That this law satisfies the basic requirement (8) is easily seen. Simply solve for the steady-state output T_∞ in response to a constant input $S = S_0$ by setting $dz/dt = 0$ in (21). The steady-state transmitter level z_∞ is

$$z_\infty = \frac{AB}{A + S_0} \tag{22}$$

and the output is

$$T_\infty = S_0 z_\infty = \frac{ABS_0}{A + S_0}. \tag{23}$$

As desired, z is a slow accumulation–depletion variable which, by (22), *decreases* as S_0 increases; yet the output T *increases* as S_0 increases because T is a joint function of the input S and the accumulation–depletion variable z. When slowly accumulating transmitter gates that possess these properties are embedded in a tonically aroused competitive anatomy (Figure 22), properties that rationalize a large body of data can be derived. When LTM rules such as (17) are clearly dsitinguished from accumulation–depletion rules such as (21), the theory is led to a pharmacological interpretation of conditioned reinforcer and motivational properties in terms of two distinct transmitter systems, whose most probable interpretation in the light of available data seemed in 1972, and still seems now in the light of a vastly expanded data base, to describe cholinergic–catecholaminergic interactions (Butcher, 1978; Friedhoff, 1975*a*,*b*).

Perhaps the single most surprising dipole property is the fact that a nonspecific arousal increment which occurs while the on-channel of a gated dipole is active can cause an antagonistic rebound. This fact was soon realized to mean that a surprising event, by triggering an arousal increment, can disconfirm and thereby rapidly extinguish the motivational support of a conditioned reinforcer. The same action of a surprising event can also enhance the rewarding effect of a reward which occurs on a trial when nonreward is expected. Then the rebound due to the reward's

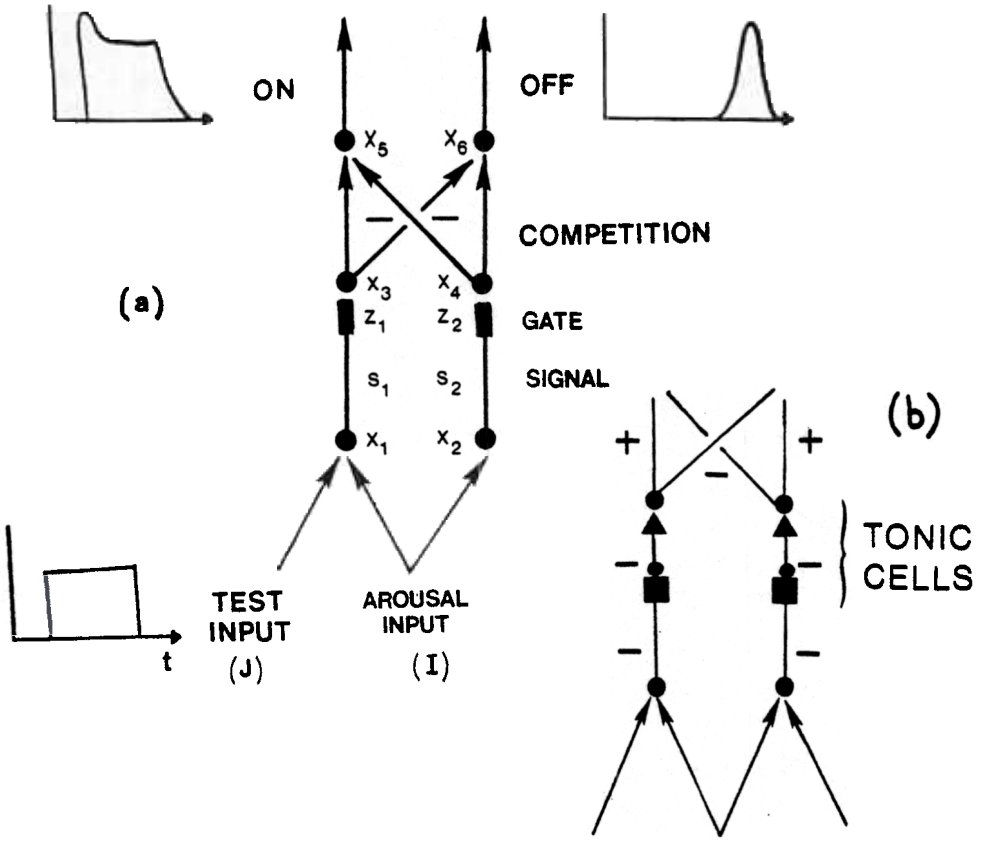

Fig. 22. Two examples of gated dipoles. In (*a*), the phasic input J and the arousal input I add in the on-channel, thereby activating the STM trace x_1. The arousal input I also perturbs the STM trace x_2 in the off-channel. Consequently, $x_1 > x_2$. Then x_1 and x_2 elicit signals $f(x_1)$ and $f(x_2)$ in their respective pathways. Since $x_1 > x_2$ also $f(x_1) > f(x_2)$. Each signal is gated (multiplied) by an excitatory transmitter z_1 or z_2 (in the square synapses) before the gated signals $f(x_1)z_1$ and $f(x_2)z_2$ activate their target cells. The STM traces x_5 and x_6 then satisfy $x_5 > x_6$. Each STM trace elicits an excitatory signal down its own pathway and an inhibitory signal to the other pathway. The net effect after competition takes places is an output from the on-channel. The text describes how a rapid offset of J triggers an antagonistic rebound that transiently excites the off-channel. In (*b*), another version of a gated dipole is depicted. Here each excitatory gating pathway is replaced by a two-stage disinhibitory pathway that is constructed from two successive inhibitory links. The cells which receive these transmitter signals are assumed to be tonic (internally and persistently activated). The net effect of an input to the two-stage disinhibitory pathway is to release its output cell from tonic inhibition and to thereby excite it.

unexpectedness can summate with the unconditional action of the reward, as in partial reward paradigms. In other words, the gated dipole machinery which produces a rebound in reaction to offset of a specific cue in classical conditioning can also produce a rebound in response to nonoccurrence of an expected event in instrumental conditioning. This property necessitates a major break with drive

reduction theory, since positive and negative rewards can hereby be manipulated by changing expectancies and environmental contingencies, even if no change in drive occurs.

34. Dipole fields in motivational processing by the hippocampal–hypothalamic axis

Our derivation has led us from the synchronization problem of classical conditioning to the notion of a gated dipole that regulates the balance between on-cell and off-cell activation in each drive representation. With these facts in hand, we can achieve a deeper understanding of how to marry the concept of a sensory-drive heterarchy between several motivational channels (Section 28) with the concept of a gated dipole within each motivational channel (Section 32). This marriage leads to the concept of a *dipole field*, a concept which is also basic in the design of cortical sensory processing areas (Grossberg, 1976*b*, 1980*a*). To reach this concept, I state postulate (J^+) in somewhat more general terms as a distinct postulate.

K. Offset of a conditioned reinforcer can have reinforcing properties of opposite sign.

This postulate implies that the dipole geometry is recurrent, or a feedback geometry. This is because the conditionable LTM pathways of conditioned reinforcing cues end *after* each dipole's transmitter gating stage, so they can learn both the on-reactions and the off-reactions of the dipole. These LTM pathways also end *before* each dipole's transmitter gating stage, so their offset can elicit an antagonistic rebound, as postulate (K) requires. In order for these pathways to end both before and after the transmitter gating stage, the dipole pathways must close upon themselves in positive feedback loops. Moreover, since the on-cells and off-cells in each dipole compete, these feedback loops are part of a recurrent on-center off-surround network. Otherwise expressed, this feedback construction endows offset of conditioned reinforcing cues with secondary conditioning properties.

The competitive feedback anatomy also enjoys several other important properties, such as its ability to defend its operating level against small input fluctuations, thereby guaranteeing a stable motivational baseline, and to control sharp motivational switching between incompatible motivational alternatives. From a general information-processing point of view, these properties are at least as basic as the secondary conditioning property of postulate (K). What is important for our present purposes is to realize that a connection exists between secondary conditioning and these other properties. Also postulate (K) was the concept which originally forced my realization that feedback competitive networks play an important role in motivational processing (Grossberg, 1972*c*), and it is an accessible property whereby to instruct students who are familiar with reinforcement but unfamiliar with cognitive processing or psychophysiology.

Several general properties of competitive feedback networks are so basic for

our present needs that I will review them here (Grossberg, 1973, 1975, 1980*a*). The feedback channels excite themselves and inhibit each other. Due to the positive feedback loops, there exists the danger that the network will amplify small noise levels into large activities and thereby flood itself with noise. Noise amplification is prevented if the feedback signals are sigmoid, or S-shaped, functions of cell activity; in particular, if spiking frequency is a sigmoid function of cell potential. The faster-than-linear growth of the sigmoid signal at small activity levels guarantees noise suppression; for example, a signal function $f(w)$ that behaves like a power aw^n ($a > 0$, $n > 1$) at small values of the activity w is faster-than-linear at these activities.

The sigmoid signal also guarantees that the network behaves like a tunable filter. In other words, there exists a *quenching threshold* (QT) that must be exceeded by a cell's activity before the positive feedback signals can effectively maintain activity in its loop. If some combination of drive, arousal, and conditioned reinforcer input exceeds the QT and wins the motivational competition, then feedback signalling can amplify the activity in the winning drive representation as it competitively suppresses other drive representations until a prescribed operating level is attained and maintained through time. I identify this maintenance property with short-term memory (STM) storage in the motivational processor. Mechanisms which maintain a proper balance between the size of the QT and of arousal inputs, drive inputs, and incentive inputs are clearly of the greatest importance for achieving proper operating characteristics within a dipole field.

Figure 23 depicts a network that marries together these several processing constraints. This is not the only possibility *in vivo*, but my discussion will attempt to be sufficiently principled to abet recognition of thematic variation across species. Figure 23 builds up a dipole field in which a sensory-drive heterarchy exists between motivational channels, gated dipoles exist within motivational channels, LTM conditioning of a conditioned reinforcer pathway is driven by positive feedback within its motivational channel, and the polyvalent cells are activated by a two-stage disinhibitory reaction to drive and arousal inputs (Figure 22*b*). In Figure 23, pathways 1 and 2 correspond to specific, but complementary, drive inputs within a single dipole. Pathways labelled 3 carry the arousal signals to this dipole. Cells 4 and 5 receive these inputs and thereupon inhibit the tonically active cells 6 and 7. All tonically active cells have open symbols; phasically active cells have closed symbols. The pathways $4 \rightarrow 6$ and $5 \rightarrow 7$ are assumed to contain slow transmitter gates (square synapses; catecholaminergic?). If input 1 exceeds input 2, then the transmitter in pathway $4 \rightarrow 6$ is depleted more than the transmitter in pathway $5 \rightarrow 7$, thereby calibrating the dipole for a possible antagonistic rebound later on.

The tonic cells 6 and 7 equally inhibit each other until input 1 exceeds input 2. Then cell 6 is inhibited more than cell 7. This has the net effect of disinhibiting polyvalent cell 8 and further inhibiting polyvalent cell 9. Due to the polyvalent nature of cells 8 and 9, these events are insufficient to fire cell 8, but if the inhibition from cell 7 to cell 9 is sufficiently large, then it can prevent cell 9 from firing at all.

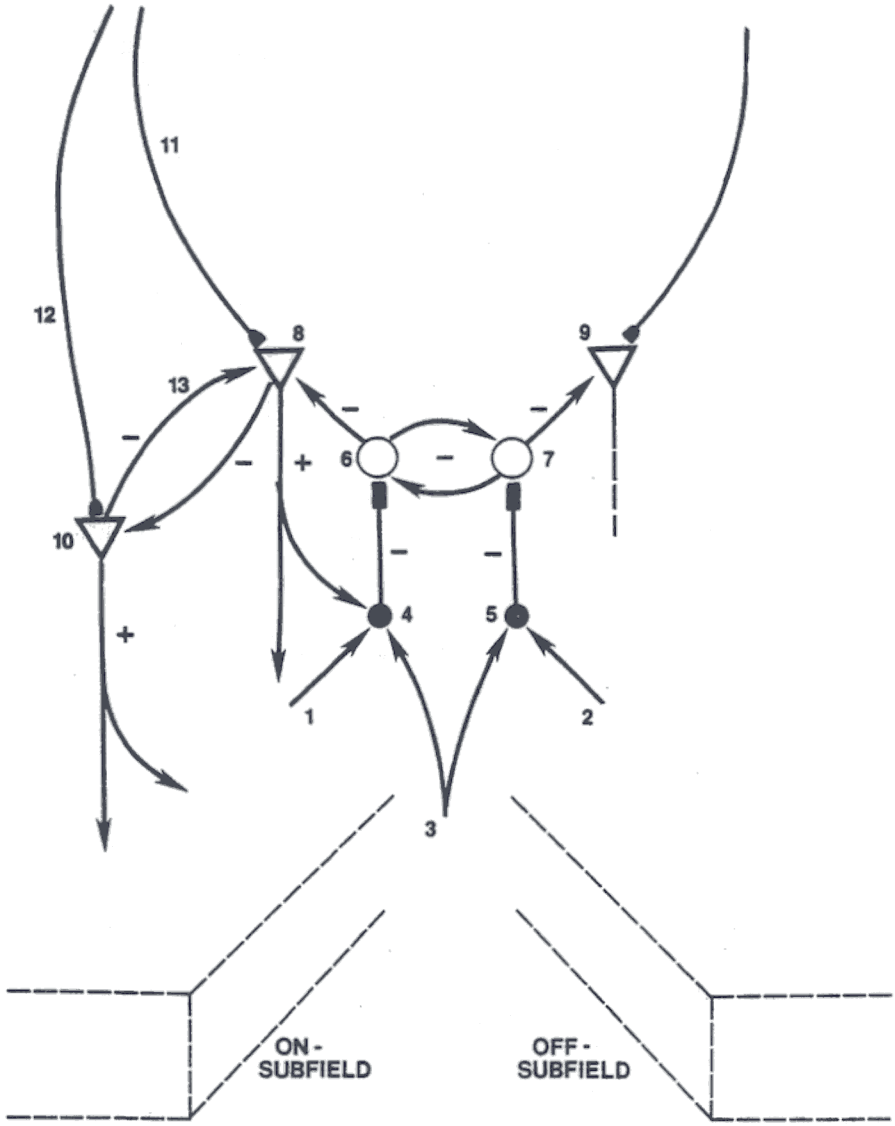

Fig. 23. A dipole field of drive representations whose conditioning, motivational, and anatomical implications are summarized in the text.

We therefore restrict our next comments to cell 8 and its competitors. Apart from the feedback between cells 6 and 7, all of our remarks to this point discuss feed-forward mechanisms. The competition between 6 and 7 could also be made feed-forward by adding some more interneurons to the network. Our next remarks are concerned with feedback properties of the network.

Cells 8 and 10 are the polyvalent cells of two different motivational channels whose drive inputs constitute on-signals, or consummatory signals. These cells receive input not only from their respective dipoles but also from conditioned reinforcing cues along LTM pathways such as 11 and 12 (cholinergic?). These conditioned reinforcer inputs combine with drive inputs at their respective polyvalent cells, which thereupon begin to generate outputs if their thresholds are exceeded. At this stage, the several polyvalent cells compete among themselves via the 'intrinsic' feedback inhibitory pathways 13 (GABA? See Siegel *et al.*, 1976), as they simultaneously try to excite themselves via positive feedback pathways. For example, cell 8 excites itself along the positive feedback pathway $8 \rightarrow 4 \rightarrow 6 \rightarrow 8$. The winner of this sensory-drive heterarchical competition has its activity stored in STM. Suppose for definiteness that the winner is cell 8. Now several things happen.

First, the positive feedback pathway $8 \rightarrow 4 \rightarrow 6 \rightarrow 8$ substantially depletes the transmitter gate in pathway $4 \rightarrow 6$. Hence *after* a motivational decision has been made, a fixed increment in arousal can cause a much larger rebound within a dipole than *before* a motivational decision has been made. This formal property can be used to experimentally test whether a transmitter gate occurs within a particular network's positive feedback loop. Second, the positive feedback amplifies the activity of cell 8 to the point where active LTM pathways (e.g., 11) abutting on cell 8 can be conditioned. This amplification of cell 8's activity might manifest itself *in vivo* as sustained bursting behavior of this polyvalent cell. Within this network, only the outcome of sensory-drive heterarchical competition can drive LTM changes. Subliminal shifts in the pattern of dipole inputs can have, at best, a small effect on LTM.

35. Some pharmacological and physiological correlates of motivational dipole fields

Some of the circuitry of Figure 23 can be used to explain the following psychopharmacological data, which are admirably reviewed by Olds (1977, pp. 59–75). Chlorpromazine and reserpine deplete amine stores. This is analogous to depleting the transmitter in the gating synapses of pathways $4 \rightarrow 6$ and $5 \rightarrow 7$. *In vivo* this manipulation depresses behavior. In Figure 23, it disinhibits the tonic inhibition of the polyvalent cells, and makes it harder for the polyvalent cells to fire. When these drugs are combined with monoamine oxidase inhibitors (MAOI), then the amines are released but they remain undegraded. This produces a lot of extra free amines. This operation would inhibit the tonic inhibitory interneurons in Figure 23. *In vivo* this manipulation abets self-stimulation, as it also would in Figure 23 by making it easier for the polyvalent cells to fire and thus for conditioned reinforcer pathways of lever press cues to get conditioned.

Amphetamines release amines and prevent their reuptake, so they act like releasers and MAOI taken together. Amphetamine can also abet self-stimulation behavior. More interestingly, amphetamines can augment slow behavior and

depress fast behavior. In small doses, amphetamine can facilitate eating, but in larger doses it can suppress eating. Such effects can be explained as an inverted U in gated dipole responsiveness when its net arousal level is parametrically increased (Grossberg, 1972*c*, 1981*a*, 1982*a*—Appendix A).

Even the highly schematic dipole field in Figure 23 requires a considerable amount of neural circuitry to carry out the competitive interactions within each dipole and between subsets of dipoles, and to deliver positive feedback signals back to the appropriate channels. A nice feature of these circuits is that simple growth rules can be suggested whereby such circuits might develop. For example, conditioned reinforcer pathways grow to polyvalent cells, dipole inhibitory interneurons grow among tonic cells, drive and arousal pathways grow to gating cells, and feedback inhibitory interneurons grow among polyvalent cells. In other words, simple chemical labels and a certain degree of symmetry in the initial network geometry can go a long way towards explaining the developmental feasibility of these circuits.

Various data suggest that these formal circuits are analogous to the neural circuits joining the hypothalamus, septum, amygdala, and hippocampus (Bridge and Hatton, 1973; De France, 1976; Haymaker *et al.*, 1969; MacLean, 1970; Olds, 1977). I have elsewhere suggested that the dipole computations are related to reciprocal hypothalamic interactions, notably between lateral and ventromedial hypothalamus; that the polyvalent cells are analogous to hippocampal pyramidal cells; and that the feedback pathways are analogous to the medial forebrain bundle (Grossberg, 1972*c*, 1975). This interpretation implies that certain hippocampal pyramidal cells are influenced by both the conditioned reinforcing properties of signals of cortical origin and the drive properties of signals of hypothalamic origin; that these pyramidal cells compete via recurrent inhibitory interneurons; that the bursting behavior of these pyramidal cells can drive conditioned changes; and that the output of these pyramidal cells is important for the transfer of STM into LTM not only directly by driving conditioned reinforcer changes but also indirectly via the effects of incentive motivational signals on the firing of cortical polyvalent cells. One might particularly note that self-stimulation is suppressed by either upstream or downstream blockade of the medial forebrain bundle (Stein, 1958). Of course the formal circuits to not presume to include all the cells that are needed to conduct business between hypothalamus and hippocampus *in vivo*. Nonetheless, the circuits help to explain a surprisingly large body of data, they illustrate how simple environmental pressures can lead to circuits of the type found *in vivo*, and they embody principles of network design which are robust enough to illuminate significantly more complex circuits.

36. Competition, normalization, and STM among sensory representations

In fact, these same principles can be used to explain a variety of perceptual, cognitive, and neocortical properties that are presently inexplicable without them. To quickly show why this is true, I ask the reader to consider the following dilemma.

Vigorous firing of polyvalent cells in Figure 23 should occur only if sufficiently large conditioned reinforcer and drive inputs combine to exceed the polyvalent thresholds. However, different sensory events can excite very different numbers of cortical feature detectors. If conditioned reinforcer inputs from a simple tone are enough to fire the polyvalent cell given a simultaneous drive input, what prevents conditioned reinforcer inputs from billions more cells from firing the cell without a simultaneous drive input? Even worse, if the latter input isn't too big, then won't the tone input be too small? Clearly the *total* input from these internal representations must be insensitive to the total number of active cells. Otherwise expressed, the total suprathreshold activity of the sensory field must be normalized, or conserved. A constraint on the decision rules whereby the sensory-drive heterarchy maintains its sensitivity to drive and reinforcer inputs hereby leads to the following postulate about the limited capacity of STM across cue representations.

L. The total activity across the internal representations of sensory cues is regulated.

Postulate (L), just like postulate (K), can be derived from more basic considerations, in this case considerations about how a sensory field can accurately register patterned data at all (Grossberg, 1980*a*, Appendices C and D). Postulate (L) has the virtues of relating a property of sensory fields to a property of motivational decisions, and of being immediately accessible.

The sensory inputs to the motivational polyvalent cells come from the first stages $\{v_{i1}^{(2)}\}$ at which STM reverberation takes place. Postulate (L) requires that interactions occur among these representations to prevent their outputs from increasing linearly with the number of excited sensory cells. These interactions must be inhibitory to keep the total output from growing without bound. These negative interactions must be feedback interactions to balance the positive feedback that maintains sensory STM. In all, the populations $v_{i1}^{(2)}$ are joined together by recurrent on-center off-surround interactions. We are now faced with a question on whose answer the life of the theory depends: Can competitive feedback interactions regulate the total activity of their network?

It is most gratifying that the answer is 'yes' if the interactions are of shunting type (which is the type that the membrane equations of neurophysiology obey) and if the feedback inhibition is of sufficiently long range across the field of populations (Grossberg, 1973, 1980*a*—Appendix D). This is, moreover, the same formal property that maintains a steady motivational operating level in the feedback competitive networks of the sensory-drive heterarchy. I call this formal property *normalization*. The need for a normalization property of one sort or another has long been recognized in psychology. Even Freud writes: 'The nervous system endeavors to keep constant something in its functional condition that may be described as the "sum of excitation"' (Freud and Breuer, 1959, p. 30). Maintaining this sum in a competitive network whose shunting interactions approximate multiplicative rules helps to explain the partial successes of probabilistic models

in studies of coding and memory. However, the probabilistic axioms do not well match the functional transformations of feedback competitive networks (Grossberg, 1978*a*). As often occurs in the history of science, an intuitively plausible property must go through several stages of mathematical explication before it can be cast in a mechanistic framework that reveals all of its implications.

One of these implications occurs in the competition among the network's external cue representations no less than in the network's sensory-drive heterarchy. A quenching threshold exists if the competitive feedback signals are chosen to be sigmoid functions of their activity levels. This QT will shortly be seen to be of major importance in attentional processing, since it helps to regulate which cues are attended and which are suppressed.

The cells $\{v_{i1}^{(2)}\}$ can now also be seen to be the on-cells of a dipole field. Just as offset of conditioned reinforcer input to a drive representation can trigger an antagonistic rebound in the complementary drive representation of its gated dipole, offset of sensory cues must be able to transiently excite off-cells whose output signals can be used to learn conditioned reactions to the offset event; for example, to learn to push a lever in response to the offset of a light. Thus the sensory processor is envisaged to be not merely a shunting competitive feedback network, but rather a pair of such networks organized as a dipole field. To drive antagonistic rebounds between the on-cells and off-cells of this field, there must exist a nonspecific arousal system gated by a slowly varying transmitter (catecholaminergic?), as well as specific sensory on-inputs and off-inputs that play the role in sensory representations which specific drive inputs play in drive representations.

The symmetry of the total network is hereby extended once again to include dipole fields in both sensory and drive representations. I view the existence of such dipole fields as a basic principle of cortical architecture, whether it be the neocortical architecture of sensory representations or the paleocortical architecture of the hippocampus and its attendant structures.

37. Attention and the persistence paradox of parallel processing

Because of this increased symmetry in the overall network geometry, another asymmetry now becomes apparent. At the drive representations, both specific internal inputs (drive inputs) and nonspecific external inputs (conditioned reinforcer inputs) combine to control which representations will be stored in STM by the feedback competition. At the sensory representations, this is not yet true. Only the specific external inputs (external cue inputs) regulate the competitive feedback at the first stages $\{v_{i1}^{(2)}\}$. The nonspecific internal inputs (incentive motivational inputs) control only the polyvalent second stages $\{v_{i2}^{(2)}\}$, which thus far have been excluded from the STM competition. Are there pressing psychological reasons, apart from symmetry considerations, which require us to include the polyvalent cells in the STM competition? Can the polyvalent cells be included in the STM feedback exchange without destroying the functional requirements that led us to distinguish the stages $\{v_{i1}^{(2)}\}$ from $\{v_{i2}^{(2)}\}$ in the first place?

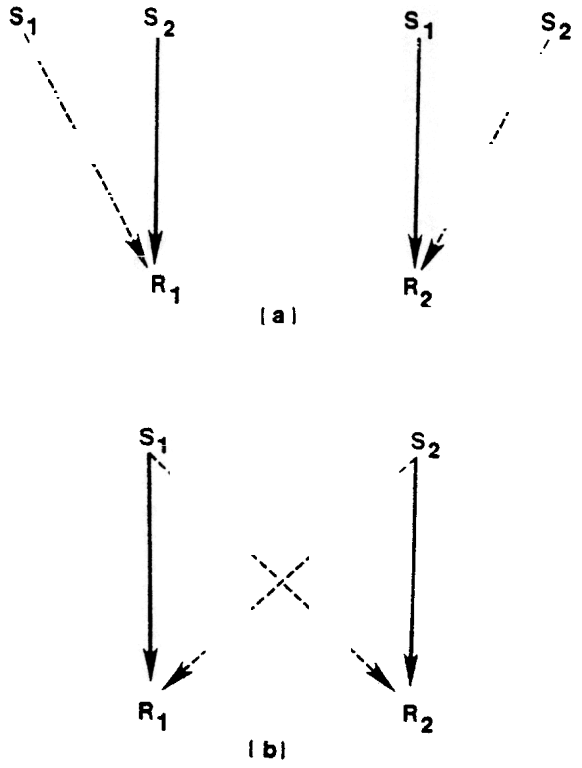

Fig. 24. Classical conditioning cannot be a passive feedforward process during real behavior. In (*a*), S_1 acts as a CS for S_2, whereas S_2 acts as a CS for S_1. In (*b*), parallel processing of S_1 and S_2, each previously conditioned to responses R_1 and R_2, would yield cross-conditioning (dotted lines).

The answer to these questions is 'yes'. The physiological requirements are illustrated by the following thought experiment, which shows how disastrous our associative links could become if classical conditioning were controlled by a feed-foward network (Grossberg, 1975, 1980*a*). In Figure 24*a* two classical conditioning experiments are depicted, one in which stimulus S_2 is the UCS for response R_2 and S_1 is its CS, and one in which S_1 is the UCS for R_1 and S_2 is its CS. What would happen if each cue S_1 and S_2 is conditioned to its own response R_1 or R_2, respectively, before a classical conditioning experiment occurs in which S_1 and S_2 are alternately scanned? This is the typical situation in real life, where we scan many cues in parallel, or intermittently, and many of these cues already have their own associations. If classical conditioning were a passive feedforward process, then cross-conditioning from S_1 to R_2 and from S_2 to R_1 would rapidly occur, as in Figure 24*b*. By contrast, we know from daily experience that the persistence of learned meanings can endure despite the fact that cues that are processed in

parallel often control incompatible responses. What feedback mechanisms of our attentional processor prevent rapid associative short-circuits from occurring?

Stated in a theoretical vacuum, this dilemma does not imply a particular mechanism, which is one reason why, despite its accessibility to all of us, it is not often even mentioned as a dilemma. At the present juncture of our theoretical construction, by contrast, the dilemma leads us to include the polyvalent cells $\{v_{i2}^{(2)}\}$ in the STM competition. Before showing why this is so, I summarize this constraint as postulate (P^4), where P^4 reminds us both of the alliterative aspects of the section title and of the pathetic consequences of not imposing it.

P^4. Persistence of learned meanings during parallel processing of motivationally imcompatible cues is possible.

Our mechanistic task is to prevent S_1 from becoming a conditioned reinforcer of R_2 and S_2 from becoming a conditioned reinforcer of R_1. Even if S_1 and S_2 are simultaneously scanned, we must prevent their first stages $v_{11}^{(2)}$ and $v_{21}^{(2)}$, respectively, from sending sustained sampling signals to the wrong drive representation. To achieve this in the present theory, at least three stages of processing are needed, two of which already exist in the network as it stands:

(1) Whe S_1 and S_2 are scanned, $v_{11}^{(2)}$ and $v_{21}^{(2)}$ send conditioned reinforcer signals to the drive representations in order to test which drive channels they control. This step answers the question: Do S_1 and S_2 control incompatible drives?

(2) The sensory-drive heterarchy determines which drive representation is stronger at each moment.

(3) Somehow $v_{i1}^{(2)} \rightarrow \mathscr{D}$ sampling signals are shut off in the weaker channel.

If property (3) can be achieved, then incentive motivational sampling will be restricted to motivationally compatible cues at each time, and attentional switching can change the class of sampling cues. How do the first stages $\{v_{i1}^{(2)}\}$ know which drive representation is stronger at any time? Feedback from drive representations to sensory representations is necessary to achieve this property. By Section 25 the incentive motivational feedback from the drive representations $\mathscr{D}$ to the second stages $\{v_{i2}^{(2)}\}$ is conditionable. Hence only these second stages which are compatible with the winning drive representation will receive strong feedback. Postulate (P^4) now suggests that when these polyvalent cells $\{v_{i2}^{(2)}\}$ fire, they send positive feedback to their respective first stages $\{v_{i1}^{(2)}\}$. In all, the incentive motivational feedback from $\mathscr{D}_i$ will excite only $v_{i1}^{(2)}$ $(i = 1,2)$. If $\mathscr{D}_1$ is the winner, the STM activity of $v_{11}^{(2)}$ will hereby be amplified. How does this amplification rapidly inhibit $v_{21}^{(2)}$ to prevent it from sampling $\mathscr{D}_1$ in a sustained fashion?

The answer is obvious due to the fact that the total activity of the first stages is normalized by the competition for STM activity which takes place among the external cue representations (Section 36). Increasing the STM activity of $v_{11}^{(2)}$ automatically decreases the STM activity of $v_{21}^{(2)}$ via the competitive feedback between these representations. This competitive feedback can, moreover, totally suppress $v_{21}^{(2)}$ if it drives its activity below the quenching threshold.

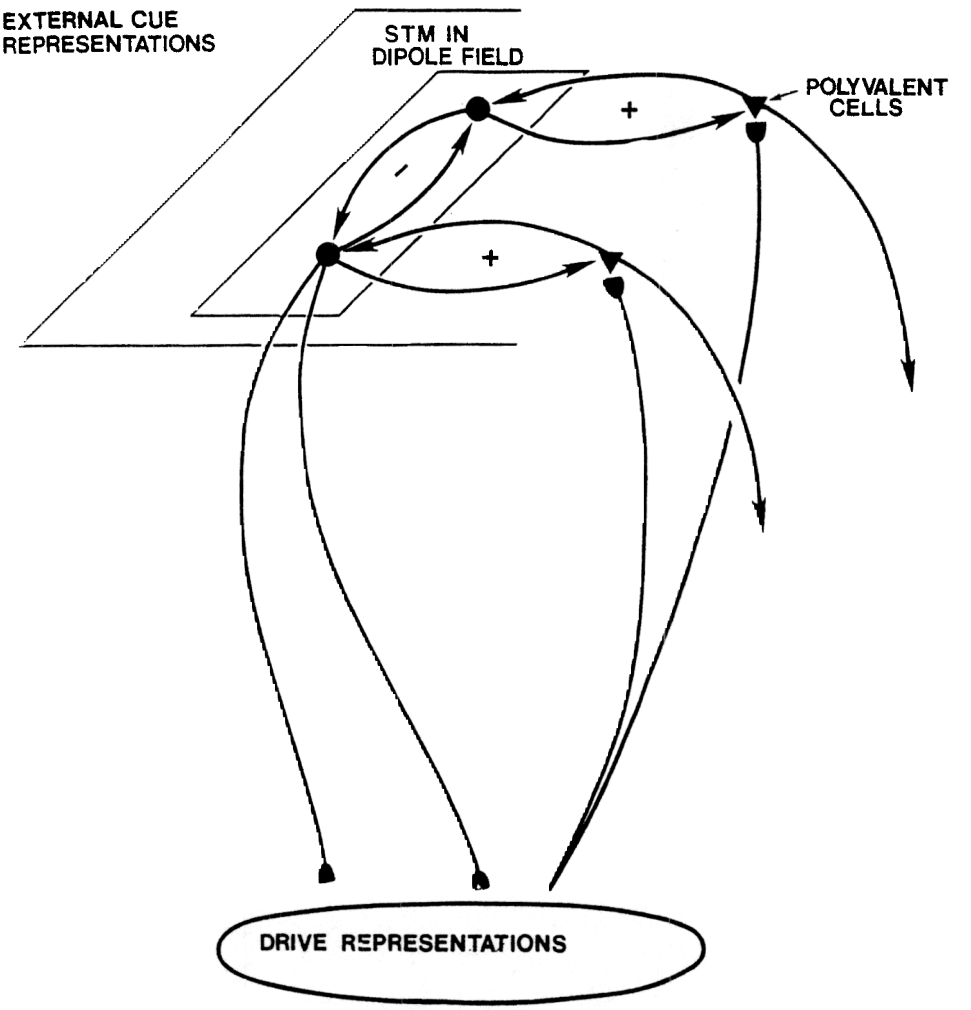

Fig. 25. Suppose that a set of nodes, or cells, in the dipole field is activated by an external scene. A pattern of STM activity across the nodes represents the scene. Each such node sends an excitatory signal to its polyvalent node, or cell. Signal size is an increasing function of STM activity. These specific signals are insufficient to fire the polyvalent cells. Sufficiently large incentive motivational signals from a drive representation must simultaneously converge upon the polyvalent cells to fire them.

The incentive motivational pathways are conditionable. A drive representation will therefore preferentially activate those polyvalent cells whose cues were paired with this drive representation in the past. The drive representation can hereby fire a subset of the polyvalent cells which are activated by the external scene. The relative rate of firing of each polyvalent cell will depend jointly on the STM activity of its trigger cues in the scene and on the relative size of its LTM trace in the conditioned reinforcer pathway. When a polyvalent cell fires, it delivers positive feedback signals to the cue-activated cells which supply it with specific STM signals. This positive feedback from polyvalent cells selectively augments the STM activities of certain cue-activated cells, which thereupon can more strongly inhibit the STM of other representations in the dipole field using the STM normalization property. The incentive motivational properties of certain cues hereby alter the set of cues to which the network pays attention. The polyvalent cells which can maintain their firing can also read out learned patterns (e.g., motor commands) to other parts of the network.

The minimal solution to the dilemma imposed by postulate (P^4) is depicted in Figure 25. In Figure 25, the polyvalent cells are part of the STM competition, but they do not fire unless they receive incentive motivational signals, so that we have not destroyed the functional constraint that forced us to distinguish the cells $\{v_{i1}^{(2)}\}$ from $\{v_{i2}^{(2)}\}$ in the first place. This construction was used in Grossberg (1975) to explain a variety of data about attention and discrimination learning. A more recent body of data are discussed using these mechanisms in Grossberg (1982*b*).

38. Sensory incentive versus motor incentive: the hippocampus as a cognitive map

To abet a fuller understanding of Figure 25, I should review a distinction from the Grossberg (1975) article which has recently been partially confirmed by data suggesting that the hippocampus is a cognitive map (O'Keefe and Nadel, 1978). In its strongest form, the O'Keefe and Nadel hypothesis claims an absolute spatial map of an animal's world in which place unit cells code where an animal is located within this map at any time. This view contrasts with data which describe the occurrence of classical conditioning at hippocampal pyramid cells using the rabbit nictitating membrane paradigm (Berger and Thompson, 1978). These latter data can be interpreted as LTM changes within conditioned reinforcer pathways at the polyvalent cells of the sensory-drive heterarchy. Various data summarized by O'Keefe and Nadel (1978) can also be interpreted using the sensory-drive heterarchy. These data probe the effects which unexpected events have on consummatory activity. Within the sensory-drive heterarchy, an unexpected event can rapidly terminate consummatory activity and release a complementary mode of orienting activity by causing antagonistic rebound of the consummatory activity's motivational source. For example, hippocampectomized rats do not orient to a novel stimulus while they are indulging in a consummatory activity, such as running towards a reward. They cannot 'shift attention during the presentation of a novel stimulus or in a mismatch situation' (O'Keefe and Nadel, 1978, p. 250). This type of effect can be qualitatively understood in the theory as it stands.

Various data concerning the manner in which animals explore an environment until they learn its spatial relationships cannot be understood unless we advance the theory further. This deficiency holds even though their notion of an absolute spatial map can be severely criticized on philosophical, no less than scientific, grounds. A weaker notion cannot, however, be so easily criticized; namely, that of a bilaterally organized motor map of approach and avoidance gradients which are built up from signal patterns that are biased by motivationally excitatory and inhibitory conditioned pathways. I shall now indicate how such a view is suggested by postulate (P^4).

The mechanism which solves postulate (P^4) works if the feedback from drive representations to sensory representations is positive. The case of drives, such as fear and frustration, which have a negative motivational sign requires further

argument. In other words, we have not yet exploited the dipole structure of the sensory-drive heterarchy.

The problem is this. If the conditioned feedback from a negative drive representations to sensory representations were negative, then it would differentially suppress activity in the corresponding sensory representations, rather than enhance their activity, in violation of postulate (P^4). Increasing the learned fearfulness of a given cue, in a fixed context of other cues, would hereby decrease the attention paid to it. This would be a most maladaptive property. Moreover, fearful cues could not overshadow or block learning in response to other cues, which is false (Kamin, 1968, 1969).

Hence a distinction must be made between channels which regulate learned persistence of negative meanings and channels which carry negative incentive motivation (Figure 26). The former channels help to focus attention on meaningful cues. Whether these cues have a positive or a negative meaning, the feedback which they control is positive. As a consequence, sensory attention can be differentially focused on these cues as they become more meaningful. The latter channels have a motor significance. They help to control approach of positive cues and avoidance of negative cues.

This distinction suggests that the output from the sensory-drive heterarchy bifurcates. One pathway carries conditionable feedback signals to sensory cortical representations. These feedback signals are analogous to the contingent negative variation (Section 25). The other pathway carries motivational signs to a field of sensory representations which lies further downstream in the network. The motivational signs differentially weight the activities of their respective cue representations, with positive signs having an excitatory effect and negative signs an inhibitory effect. These differentially weighted cue representations then project to a bilaterally organized motor map, wherein the bilateral asymmetry of the map's activity pattern at any time controls the network's approach and avoidance tendencies at this time. I suggest that the O'Keefe and Nadel (1978) data on place learning are probing properties of this latter pathway.

39. Expectancy matching and attentional reset: unblocking and dishabituation

The introduction of dipole field structure into the sensory representations propels our motivational theory deeper into the realm of perceptual and cognitive theory, and indeed no mature motivational theory can entirely avoid discussion of cognitive influences. For the sake of completeness, I will briefly review how the mismatch of feedback expectancies with feedforward data patterns can reset a dipole field by triggering an increment in its nonspecific arousal level. I will also indicate how the same arousal increment which causes antagonistic rebounds in dipoles that have previously been very active can simultaneously enhance the onreactions of dipoles that have previously been only weakly active and of the novel cues that caused the arousal increment. In other words, expectancy mismatch due to an unexpected event can disconfirm previously active representations as it

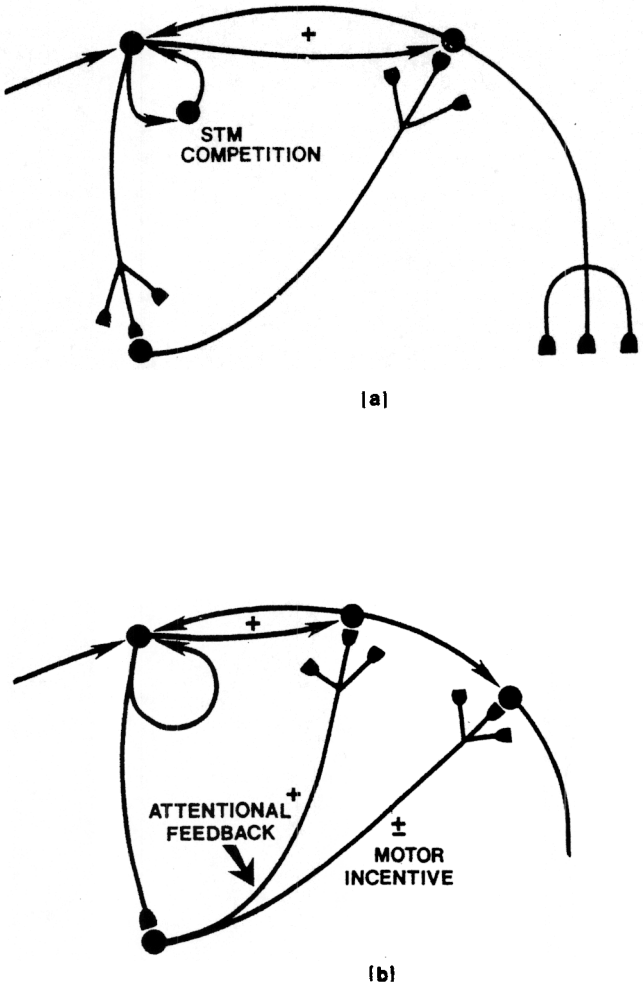

Fig. 26. In (*a*), positive feedback from $v_{12}^{(2)}$ to $v_{11}^{(2)}$ can be used to selectively enhance sensory representations which are compatible with a given incentive, and thereby to indirectly overshadow other sensory representations via the STM competition. In (*b*), the positive attentional feedback pathway is distinguished from the positive or negative pathway which assigns emotional signs to a cue-modulated motor map.

enhances previously suppressed and newly activated representations. I will not discuss how the internal representations which release the feedback expectancies are encoded, and how once encoded they can learn the proper feedback expectancy. For these discussions, the reader might consult Grossberg (1972*a*, 1976*a*,*b*, 1978*b*, 1980*a*).

The expectancy-matching mechanism is an automatic property of certain feedback competitive networks. When such a network is designed to suppress uniform

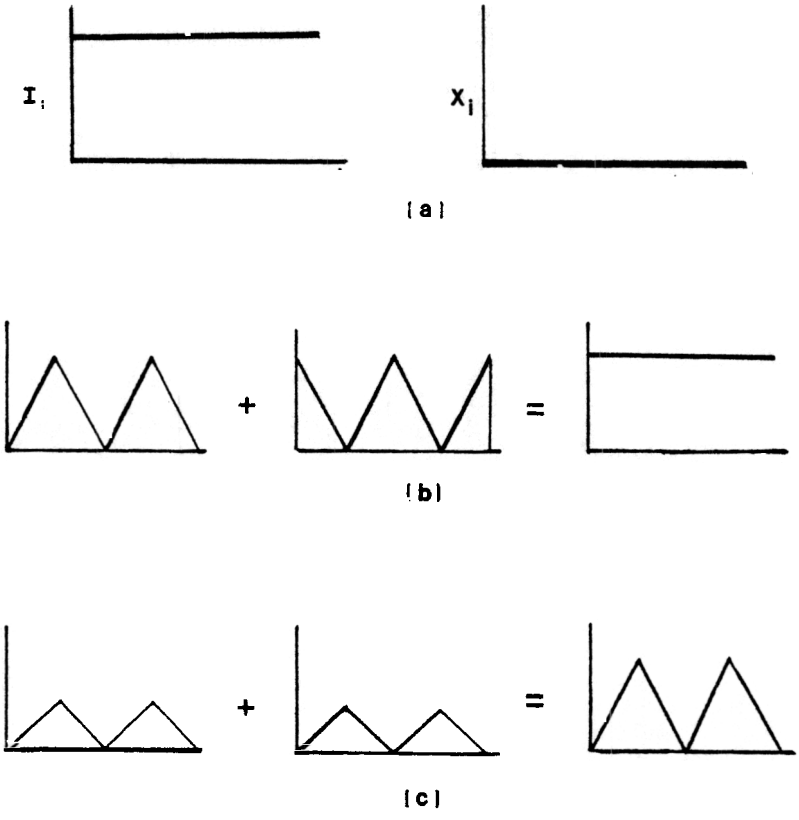

Fig. 27. In (*a*), a uniform pattern of inputs I_i is transformed into a zero pattern of activities X_i. In (*b*), two mismatched patterns add to generate an approximately uniform total input pattern, which will be suppressed by the mechanism of (*a*). In (*c*), two matched patterns add to yield a total input pattern that can elicit more vigorous activation than either input pattern taken separately.

patterns (Figure 27*a*), a sum of two mismatched input patterns will also be suppressed, since the peaks and troughs of one pattern will tend to correspond to the troughs and peaks of the other pattern, respectively (Figure 27*b*). By contrast, such a network reacts to two matched patterns by amplifying its activities (Figure 27*c*). These properties are due to automatic gain control by the inhibitory feedback signals. They are reviewed in Grossberg (1980*a*, Appendix C). In the present theory, these feedback competitive networks are in the on-cell subfield and off-cell subfield of each dipole field.

The expectancy mismatch mechanism is illustrated in Figure 28. In Figure 28, an afferent data pattern elicits activity across a competitive network $\mathscr{F}^{(1)}$ as it also activates a nonspecific arousal source $\mathscr{A}$. The activity within $\mathscr{F}^{(1)}$ rapidly inhibits the arousal source. The pattern across $\mathscr{F}^{(1)}$ then elicits signals to the next

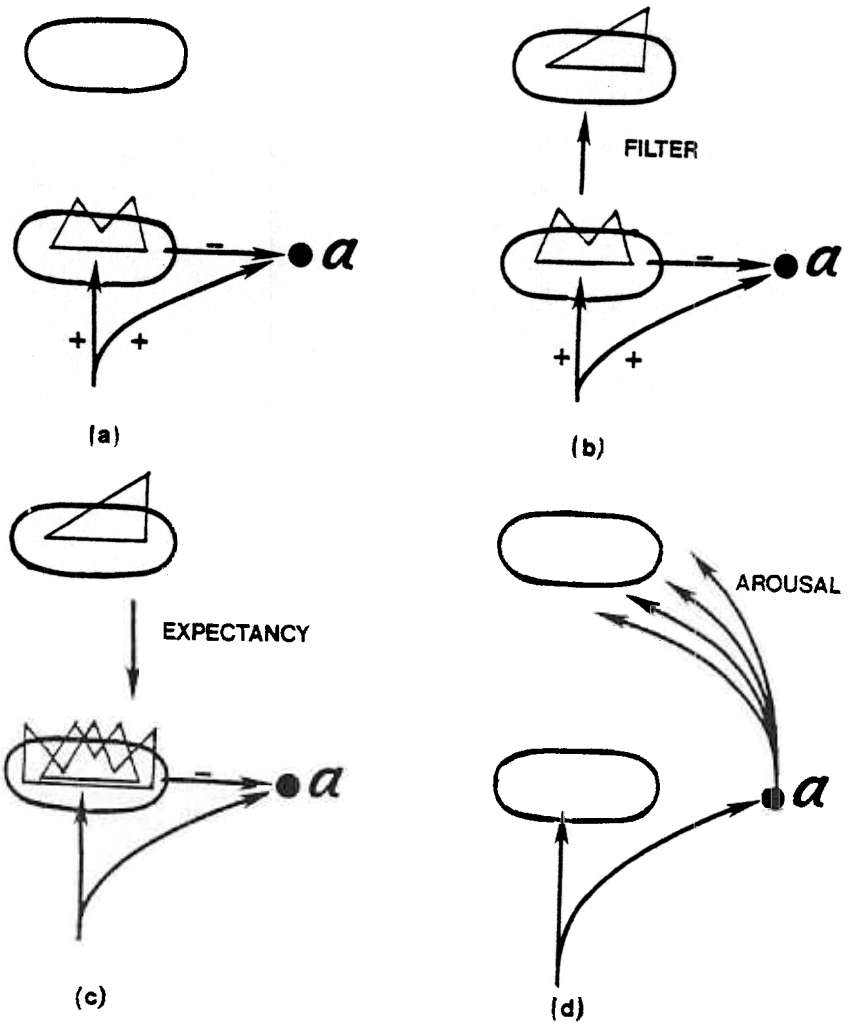

Fig. 28. In (*a*), afferent data elicit activity across $\mathscr{F}^{(1)}$ and an input to the arousal source $\mathscr{A}$ that is inhibited by $\mathscr{F}^{(1)}$. In (*b*), the pattern at $\mathscr{F}^{(1)}$ maintains inhibition of $\mathscr{A}$ as it is filtered and activates $\mathscr{F}^{(2)}$. In (*c*), the feedback expectancy from $\mathscr{F}^{(2)}$ is matched against the pattern at $\mathscr{F}^{(1)}$. In (*d*), mismatch attenuates activity across $\mathscr{F}^{(1)}$ and thereby disinhibits $\mathscr{A}$, which releases a nonspecific arousal signal to $\mathscr{F}^{(2)}$.

stage $\mathscr{F}^{(2)}$ of network processing. These signals act like an adaptive filter. The patterned output of this filter is contrast enhanced and normalized as it is read into STM at $\mathscr{F}^{(2)}$. This STM pattern, in turn, reads out a learned feedback template or expectancy to $\mathscr{F}^{(1)}$. If the expectancy mismatches the afferent data pattern at $\mathscr{F}^{(1)}$, then activity across $\mathscr{F}^{(1)}$ is inhibited, whereupon the arousal source $\mathscr{A}$ is

disinhibited. A pulse of arousal is hereby released and rapidly resets the filtered chunks which were active in STM across $\mathscr{F}^{(2)}$. The two successive stages $\mathscr{F}^{(1)}$ and $\mathscr{F}^{(2)}$ of pattern processing are part of the network's *attentional* subsystem. The arousal channel is part of the network's *orienting* subsystem. The inhibitory link between the two subsystems illustrates the complementary relationship between attentional and orienting reactions (Grossberg, 1975, 1982*a*,*b*).

Why the arousal pulse can rebound very active dipoles and enhance the on-reactions of weakly active dipoles can only be understood by mathematically analyzing gated dipole dynamics, as in Grossberg (1982*a*, Appendix A). Such an analysis indicates how a suitably chosen sigmoid signal function can determine the minimal arousal increment $g(I,J)$ which can rebound a dipole whose net arousal level is I and whose specific input in its on-channel exceeds I by J. One finds that

$$g(I,J) = \frac{A - I(I+J) + (A+I^2)^{1/2}[A+(I+J)^2]^{1/2}}{2I+J}.$$

Function $g(I,J)$ is a decreasing function of J. In other words, a fixed arousal increment ΔI can more easily rebound dipoles which have large prior on-reactions J than dipoles with small prior on-reactions J. If $\Delta I > g(I,J)$, a rebound occurs. If $\Delta I < g(I,J)$, an enhanced on-reaction occurs. By (24),

$$g(I,0) = AI^{-1}$$

whereas

$$g(I,\infty) = (A+I^2)^{1/2} - I. \tag{26}$$

Consequently $g(I,\infty)$ can equal any fraction $1/n < \frac{1}{2}$ of $g(I,0)$, viz.,

$$g(I,\infty) = \frac{1}{n}\, g(I,0), \tag{27}$$

if

$$I = \sqrt{\frac{A}{n(n-2)}}. \tag{28}$$

In other words, $g(I,J)$) can decrease to an arbitrarily small fraction of its maximal value $g(I,0)$ as J increases. This means that the mechanism whereby a fixed ΔI can differentially trigger antagonistic rebound or an enhanced on-reaction is robust. Interestingly, the size of I determines how big this effect can be, since (28) shows that the relative effect increases as I decreases. One can also readily check that the on-reaction of a dipole is also enhanced if a specific input J accompanies an increment ΔI in arousal just so long as the total input remains in the faster-than-linear part of the signal function $f(w)$. This property helps to explain the enhanced STM storage of a novel event.

The remainder of this section indicates that the attentional and orienting

mechanisms depicted in Figure 28 need not be housed in physically disjoint networks. They can be diffusely interspersed among one another, much as the visual cortical projections of retinal X-cells and Y-cells mutually interpenetrate (Breitmeyer, 1980; Robson, 1975). I offer this example not to exhaust the possibilities, but to stimulate further study of this question.

Figure 29 depicts one possible network realization of this type. The dotted lines 1 described a nonspecific arousal pathway which terminates at the specific

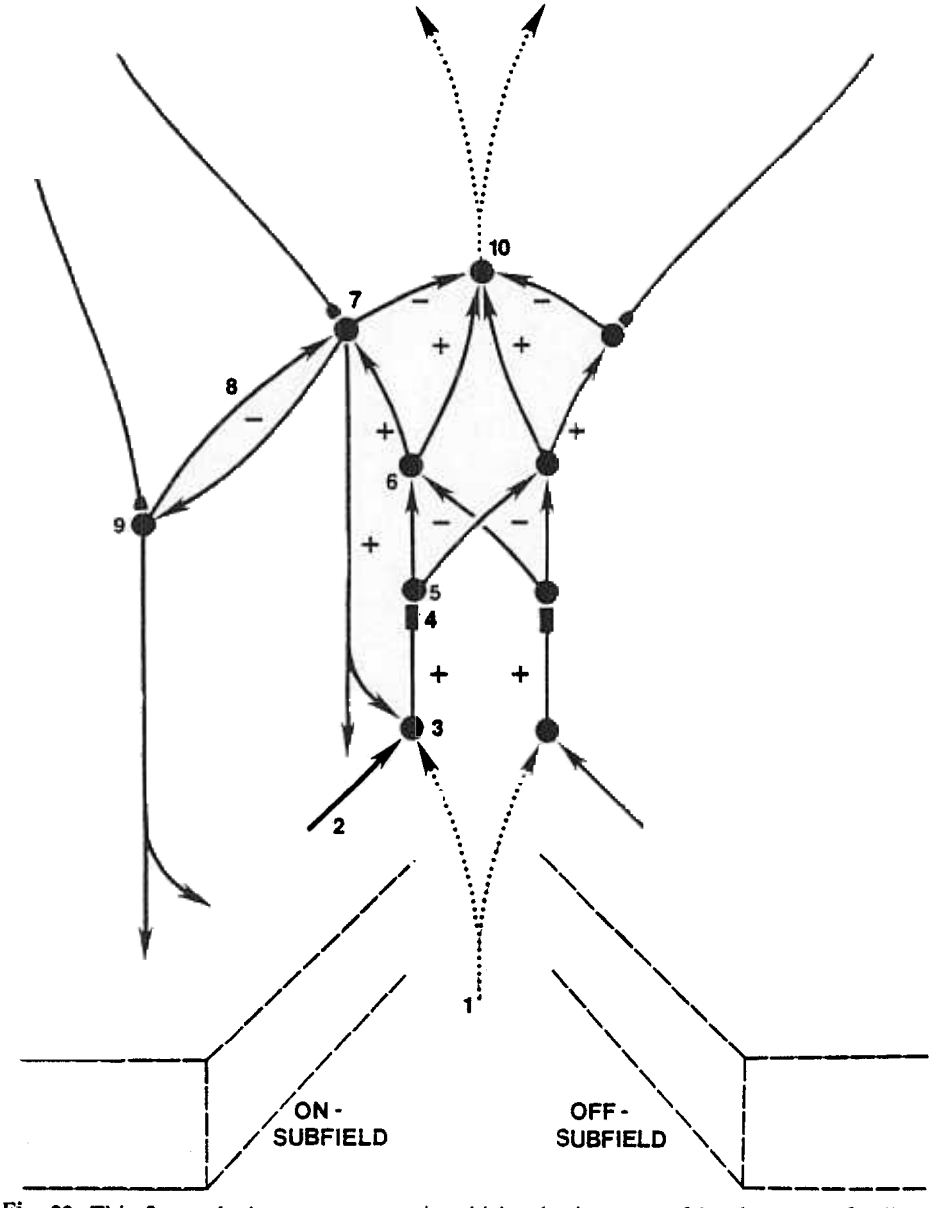

Fig. 29. This figure depicts an anatomy in which selective reset of local groups of cells can be achieved until the entire tissue can achieve resonance. The text describes the network's operation step by step.

on-cell and off-cell of the depicted dipole. Pathway 2 is the specific on-cell input to the on-cell 3. The remainder of the figure through cell 6 depicts a standard dipole geometry. I have made the transmitter gate 4 excitatory to avoid making the figure unnecessarily complicated. Stage 6 sends inputs into the feedback anatomy of the network. The positive feedback loop $6 \leftrightarrow 7$ is the recurrent on-center corresponding to the dipole's on-cell axis. Looking along the axis of the on-subfield, one sees negative feedback pathways, such as 8, to the on-channels of other dipoles, such as cell 9. This competitive feedback nework will compute pattern matches.

The new additions to the network are the arousal cells, such as 10, which are interspersed among the feedback interneurons. All of the interneurons such as 7, 9, and 10 would be considered part of 'intrinsic' feedback circuits if they were observed in cortical tissue. Cell 10 receives a positive input from cell 6 just as cell 7 does. What is different is that cell 7 inhibits cell 10, as do all the continguous dipole-related cells. The inhibitory link $7 \rightarrow 10$ prevents cell 10 from firing when cell 7 is on, just as in Figure 23.

A pattern of feedback expectancy signals is also delivered to cells such as 7 and 9. If this feedback pattern matches the feedforward data pattern across the dipole on-cells, then inhibition of cell 10 is maintained as the network resonates the feedback template pattern. If a mismatch occurs, then cell 7 is inhibited, but cell 6 is not. In other words, interneurons such as 7 and 9 form a matching interface. Cell 10 is hereby disinhibited and elicits an arousal burst to the next processing stage. An interesting feature of this construction is that the arousal can be distributed among local groups of cells. One need not reset the whole tissue due to a mismatch in just one of its spatial channels. Local readjustments of templates with data can occur until the entire tissue can achieve resonance. Otherwise expressed, these local arousal sources can cause readjustments in system coding and tuning on a channel-by-channel basis, and can thereby bias the processing in contiguous channels until the entire tissue achieves resonance.

40. Concluding remarks

This article illustrates how a real-time analysis of an individual's adaptive moment-by-moment behavior in prescribed environments can disclose network principles and mechanisms that admit a physiological and pharmacological interpretation. Each of these principles leads to a class of mathematical design problems, each of these problems can be solved, and in the light of the solutions, an interdisciplinary restructuring and unification of the data is implied in terms of design principles and mechanisms rather than the vicissitudes of experimental methodology or historical accident. This article has focused on the designs which are forced by the synchronization problem and the persistence paradox, and has used these designs to explicate the ideas contained in various classical theories and data. Related articles in this series (Grossberg, 1981*a*, 1982*a*,*b*) have used similar ideas to explain a wide variety of recent interdisciplinary data and have suggested interdisciplinary experiments to independently cross-check these ideas.

References

Anderson, P. and Eccles, J. C. (1962). Inhibitory phasing of neuronal discharge. *Nature* **196**, 645–647.

Annau, Z. and Kamin, L. J. (1961). The conditioned emotional response as a function of intensity of the US. *J. Comp. Physiol. Psych.* **54**, 428–432.

Beach, F. A. (1956). Characteristics of masculine 'sex drive'. *In: Nebraska Symposium on Motivation* (Jones, M. R., Ed.). University of Nebraska Press, Lincoln.

Berger, T. W. and Thompson, R. F. (1978). Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response. I. The hippocampus. *Brain Res.* **145**, 323–346.

Bindra, D. (1968). Neurophysiological interpretation of the effects of drive and incentive-motivation on general activity and instrumental behavior. *Psychol. Rev.* **75**, 1–22.

Bitterman, M. E. and Mackintosh, N. J. (1969). Habit-reversal and probability learning: Rats, birds, and fish. *In: Animal Discrimination Learning* (Gilbert, R. M. and Sutherland, N. S., Eds.). Academic Press, New York.

Boe, E. E. (1966). Effect of punishment duration and intensity on the extinction of an instrumental response. *J. Exp. Psych.* **72**, 125–131.

Bolles, R. C. (1967). *Theory of Motivation.* Harper and Row, New York.

Boren, J. J., Sidman, M. and Herrnstein, R. J. (1959). Avoidance, escape and extinction as functions of shock intensity. *J. Comp. Physiol. Psych.* **52**, 420–425.

Borozci, G., Storms, L. H. and Broen, W. E. (1964). Response suppression and recovery of responding at different deprivation levels as functions of intensity and duration of punishment. *J. Comp. Physiol. Psych.* **58**, 456–459.

Breitmeyer, B. G. (1980). Unmasking visual masking: A look at the 'why' behind the veil of the 'how'. *Psychol. Rev.* **87**, 52–69.

Bridge, J. G. and Hatton, G. I. (1973). Septal unit activity in response to alterations in bleed volume and osmotic pressure. *Physiol. Behav.* **10**, 769–774.

Brush, F. R. (1971). *Aversive Conditioning and Learning.* Academic Press, New York.

Butcher, L. L. (Ed.) (1978). *Cholinergic–monoaminergic interactions in the brain.* Academic Press, New York.

Campbell, B. A. and Kraeling, D. (1953). Response strength as a function of drive level and amount of drive reduction. *J. Exp. Psych.* **45**, 97–101.

Campbell, B. A. and Church, R. M. (1969). *Punishment and Aversive Behavior.* Appleton–Century–Crofts, New York.

Capaldi, E. J. (1971). Memory and learning: a sequential viewpoint. *In: Animal Memory* (Honig, W. K. and James, P. H. R., Eds.). Academic Press, New York.

Church, R. M., Raymond, G. A. and Beauchamp, R. D. (1967). Response suppression as a function of intensity and duration of punishment. *J. Comp. Physiol. Psych.* **63**, 39–44.

Cotton, J. W. (1953). Running time as a function of amount of food deprivation. *J. Exp. Psych.* **46**, 188–198.

Crespi, L. P. (1942). Quantitative variation of incentive and performance in the white rat. *Am. J. Psych.* **55**, 467–517.

Crespi, L. P. (1944), Amount of reinforcement and level of performance. *Psychol. Rev.* **51**, 341–357.

Crosby, E. C., Humphrey, T. and Lauer, E. W. (1962). *Correlative Anatomy of the Nervous System.* Macmillan, New York.

D'Amato, M. R., Fazzaro, J. and Etkin, M. (1967). Discriminated bar-press avoidance maintenance and extinction in rats as a function of shock intensity. *J. Comp. Physiol. Psych.* **63**, 351–354.

Deese, J. and Carpenter, J. A. (1951). Drive level and reinforcement. *J. Exp. Psych.* **42**, 236–238.

De France, J. F. (1976). *The Septal Nuclei.* Plenum Press, New York.

Denny, M. R. (1971). Relaxation theory and experiments. *In: Aversive Conditioning and Learning* (Brush, F. R., Ed.). Academic Press, New York.

Dickinson, A., Hall, G. and Mackintosh, N. J. (1976). Surprise and the attenuation of blocking. *J. Exp. Psych.: Animal Behavior Processes* **2**, 213–222.

Dunham, P. J. (1971). Punishment: method and theory. *Psychol. Rev.* **78**, 58–70.

Estes, W. K. (1958). Stimulus–response theory of drive. *In: Nebraska Symposium on Motivation* (Jones, M. R., Ed.). University of Nebraska Press, Lincoln.

Estes, W. K. (1969). Outline of a theory of punishment. *In: Punishment and Aversive Behavior* (Campbell, B. A. and Church, R. M., Eds.). Appleton–Century–Crofts, New York.

Estes, W. K. and Skinner, B. F. (1941). Some quantitative properties of anxiety. *J. Exp. Psych.* **29**, 390–400.

Freud, S. and Breuer, J. (1959). On the theory of hysterical attacks. *In: Collected Papers*, Volume V. Basic Books, New York.

Friedhoff, A. J. (Ed.) (1975*a*). *Catecholamines and Behavior, Vol. I: Basic Neurobiology*. Plenum Press, New York.

Friedhoff, A. J. (Ed.) (1975*b*). *Catecholamines and behavior, Vol. II: Neuropsychopharmacology*. Plenum Press, New York.

Gabriel, M., Foster, K., Orona, E., Saltwick, S. E. and Stanton, M. (1980). Neuronal activity of cingulate cortex, anteroventral thalamus and hippocampal formation in discrimination conditioning: Encoding and extraction of the significance of conditional stimuli. *Prog. Psychobiol. Physiol. Psych.* **9**, 125–231.

Gilbert, R. M. and Sutherland, N. S. (1969). *Animal Discrimination Learning*. Academic Press, New York.

Goldman-Rakic, P. S. (1981). Prenatal formation of cortical input and development of cytoarchitectonic compartments in the neostriatum of the rhesus monkey. *J. Neurosci.* **1**, 721–735.

Grossberg, S. (1964). *The Theory of Embedding Fields with Applications to Psychology and Neurophysiology*. Rockefeller Institute for Medical Research, New York.

Grossberg, S. (1968). Some physiological and biochemical consequences of psychological postulates. *Proc. Natl. Acad. Sci. U.S.A.* **60**, 758–765.

Grossberg, S. (1969*a*). On learning and energy–entropy dependence in recurrent and nonrecurrent signed networks. *J. Stat. Phys.* **1**, 319–350.

Grossberg, S. (1969*b*). On the serial learning of lists. *Math. Biosci.* **4**, 201–253.

Grossberg, S. (1969*c*). On the production and release of chemical transmitters and related topics in cellular control. *J. Theor. Biol.* **22**, 325–364.

Grossberg, S. (1970*a*). Some networks that can learn, remember, and reproduce any number of complicated space–time patterns, II. *Stud. Appl. Math.* **49**, 135–166.

Grossberg, S. (1970*b*). Neural pattern discrimination. *J. Theor. Biol.*, **27**, 291–337.

Grossberg, S. (1971*a*). On the dynamics of operant conditioning. *J. Theor. Biol.*, **33**, 225–255.

Grossberg, S. (1971*b*). Pavlovian pattern learning by nonlinear neural networks. *Proc. Natl. Acad. Sci. U.S.A.* **68**, 828–831.

Grossberg, S. and Pepe, J. (1971). Spiking threshold and overarousal effects on serial learning. *J. Stat. Phys.* **3**, 95–125.

Grossberg, S. (1972*a*). Pattern learning by functional-differential neural networks with arbitrary path weights. *In: Delay and Functional-differential Equations and their Applications* (Schmitt, K., Ed.). Academic Press, New York.

Grossberg, S. (1972*b*). A neural theory of punishment and avoidance, I. Qualitative theory. *Math. Biosci.* **15**, 39–67.

Grossberg, S. (1972*c*). A neural theory of punishment and avoidance, II. Quantitative theory. *Math. Biosci.* **15**, 253–285.

Grossberg, S. (1973). Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Stud. Appl. Math.* **52**, 217–257.

Grossberg, S. (1974). Classical and instrumental learning by neural networks. *In: Progress in Theoretical Biology*, Vol. 3 (Rosen, R. and Snell, F., Eds.). Academic Press, New York.

Grossberg, S. (1975). A neural model of attention, reinforcement, and discrimination learning. *Int. Rev. Neurobiol.* **18**, 263–327.

Grossberg, S. (1976*a*). Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors. *Biol. Cybernetics* **23**, 121–134.

Grossberg, S. (1976*b*). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biol. Cybernetics* **23**, 187–202.

Grossberg, S. (1978*a*). Behavioral contrast in short-term memory: Serial binary memory models or parallel continuous memory models? *J. Math. Psych.* **17**, 199–219.

Grossberg, S. (1978*b*). A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and plans. *In: Progress in Theoretical Biology*, Vol. 5 (Rosen, R. and Snell, F. (Eds.). Academic Press, New York.

Grossberg, S. (1978*c*). Competition, decision, and consensus. *J. Math. Anal. Appl.* **66**, 470–493.

Grossberg, S. (1980*a*). How does a brain build a cognitive code? *Psychol. Rev.* **1**, 1–51.

Grossberg, S. (1980*b*). Biological Competition: Decision rules, pattern formation, and oscillations. *Proc. Natl. Acad. Sci.* **77**, 2338–2342.

Grossberg, S. (1981*a*). Psychophysiological substrates of schedule interactions and behavioral contrast. *In: Mathematical Psychology and Psychophysiology* (Grossberg, S., Ed.). American Mathematical Society, Providence.

Grossberg, S. (1981*b*). Adaptive resonance in development, perception, and cognition. *In: Mathematical Psychology and Psychophysiology* (Grossberg, S., Ed.). American Mathematical Society, Providence.

Grossberg, S. (1982*a*). Some psychophysiological and pharmacological correlates of a developmental cognitive, and motivational theory. *In: Cognition and Brain Activity* (Cohen, J., Karrer, R. and Tueting, P., Eds.). New York Academy of Sciences, New York.

Grossberg, S. (1982*b*). The processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychol. Rev.* (in press).

Grossberg, S. (1982*c*). *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control.* Reidel, Boston.

Grossman, S. P. (1967). *A Textbook of Physiological Psychology.* Wiley, New York.

Hall, G. and Pearce, J. M. (1979). Latent inhibition of a CS during CS–US pairings. *J. Exp. Psych.: Animal Behavior Processes* **5**, 31–42.

Haymaker, W., Anderson, E., and Nauta, W. J. H. (1969). *The Hypothalamus.* C. C. Thomas, Springfield.

Hebb, D. O. (1955). Drives and the CNS (conceptual nervous system). *Psychol. Rev.* **62**, 243–254.

Hinson, J. M. and Staddon, J. E. R. (1978). Behavioral competition: A mechanism for schedule interactions. *Science* **202**, 432–434.

Honig, W. K. (1966). *Operant behavior: Areas of Research and Application.* Appleton–Century–Crofts, New York.

Hoyle, G. (1977). *Identified Neurons and Behavior of Arthropods.* Plenum Press, New York.

Hubel, D. H. and Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proc. Roy. Soc. London* (B) **198**, 1–59.

Huff, F. W., Piantanida, T. P., and Morris, G. L. (1967). Free operant avoidance responding as a function of serially presented variations in UCS intensity. *Psychon. Sci.* **8**, 111–112.

Hull, C. L. (1943). *Principles of Behavior.* Appleton, New York.

Hull, C. L. (1952). *A Behavior System.* Yale University Press, New Haven.

Irwin, D. A., Rebert, C. S., McAdam, D. W. and Knott, J. R. (1966). Slow potential change (CNV) in the human EEG as a function of motivational variables. *Electroenceph. Clin. Neurophysiol.* **21**, 412–413.

Janik, A. and Toulmin, S. (1973). *Wittgenstein's Vienna.* Simon and Schuster, New York.

John, E. R. (1966). Neural processes during learning. *In: Frontiers in Physiological Psychology* (Russell, R. W., Ed.). Academic Press, New York.

John, E. R. (1967). *Mechanisms of Memory.* Academic Press, New York.
John, E. R. and Morgades, P. P. (1969). Neural correlates of conditioned responses studied with multiple chronically implanted moving electrodes. *Exp. Neurol.* **23**, 412–425.
Johnson, J. L. and Church, R. M. (1965). Effects of shock intensity on nondiscriminative avoidance learning of rats in a shuttlebox. *Psychon. Sci.* **3**, 497–498.
Kamin, L. J. (1968). 'Attention-like' processes in classical conditioning. *In: Miama Symposium on the Prediction of Behavior: Aversive Stimulation* (Jones, M. R., Ed.). University of Miami Press, Miami.
Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. *In: Punishment and Aversive Behavior* (Campbell, B. A. and Church, R. M., Eds.). Appleton–Century–Crofts, New York.
Kamin, L. J., Brimer, C. J. and Black, A. H. (1963). Conditioned suppression as a monitor of fear in the course of avoidance training. *J. Comp. Physiol. Psych.* **56**, 497–501.
Kandel, E. R. (1976). *Cellular Basis of Behavior: An Introduction to Behavioral Neurobiology.* W. H. Freeman, San Francisco.
Keehn, J. D. (1963). Effect of shock duration on Sidman avoidance response rates. *Psychol. Repts.*, **13**, 852.
Kelleher, R. J. (1966). Chaining and conditioned reinforcement. *In: Operant Behavior: Areas of Research and Application* (Honig, W. K., Ed.). Appleton–Century–Crofts, New York.
Kelso, J. A. S., Southard, D. L. and Goodman, D. (1979). On the nature of human interlimb coordination. *Science* **203**, 1029–1031.
Kendler, H. H. (1946). The influence of simultaneous hunger and thirst drives upon the learning of two opposed spatial responses of the white rat. *J. Exp. Psych.* **36**, 212–220.
Kendler, H. H. and Levine, S. (1951). Studies of the effect of change of drive: I. From hunger to thurst in a T-maze. *J. Exp. Psych.* **41**, 429–436.
Kendler, H. H., Levine, S., Altchek, E. and Peters, H. (1952). Studies of the effect of change of drive: II. From hunger to different intensities of thirst drive in a T-maze. *J. Exp. Psych.* **44**, 1–4.
Livingston, R. B. (1967). Brain mechanisms in conditioning and learning. *In: Neurosciences Research Symposium Summaries*, Vol. 2 (Schmitt, F. O., Melnechuk, T., Quartan, G. C. and Adelman, G., Eds.). M.I.T. Press, Cambridge.
Logan, F. A. (1969). The negative incentive value of punishment. *In: Punishment and Aversive Behavior* (Campbell, B. A. and Church, R. M., Eds.). Appleton–Century–Crofts, New York.
Low, M. D., Borda, R. P., Frost, J. D. and Kellaway, P. (1966). Surface negative slow potential shift associated with conditioning in man. *Neurology* **16**, 711–782.
Macchi, G. and Rinvik, E. (1976). Thalamo–telencephalic circuits: A neuroanatomical survey. *In: Handbook of Electroencephalography and Clinical Neurophysiology*, Vol. 2, Pt. 1 (Remond, A., Ed.). Elsevier, Amsterdam.
Mackintosh, N. J. (1971). An analysis of overshadowing and blocking. *Quart. J. Exp. Psych.* **23**, 118–125.
Mackintosh, N. J. (1974). *The Psychology of Animal Learning.* Academic Press, New York.
Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychol Rev.* **82**, 276–298.
Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning and Behavior* **4**, 186–192.
Mackintosh, N. J., Bygrave, D. J. and Picton, B. M. B. (1977). Locus of the effect of a surprising reinforcer in the attenuation of blocking. *Quart. J. Exp. Psych.* **29**, 327–336.
Mackintosh, N. J. and Reese, B. (1979). One-trial overshadowing. *Quart. J. Exp. Psych.* **31**, 519–526.
MacLean, P. D. (1970). The limbic brain in relation to the psychoses. *In: Physiological Correlates of Emotion* (Black, P., Ed.). Academic Press, New York.

Martin, L. K. and Reiss, D. (1969). Effects of US intensity during previous discrete delay conditioning on conditioned acceleration during avoidance conditioning. *J. Comp. Physiol. Psych.* **69**, 196–200.

Masterson, F. A. (1970). Is termination of a warning signal an effective reward for the rat? *J. Comp. Physiol. Psych.* **72**, 471–475.

McAdam, D. W. (1969). Increases in CNS excitability during negative cortical slow potentials in man. *Electroenceph. Clin. Neurophysiol.* **26**, 216–219.

Miller, N. E. (1963). Some reflections on the law of effect produce a new alternative to drive reduction. *In: Nebraska Symposium on Motivation* (Jones, M. R., Ed.). University of Nebraska Press, Lincoln.

Mogenson, G. J., Mullin, A. D. and Clark, E. A. (1965). Effects of delayed secondary reinforcement and response requirements on avoidance learning. *Canadian J. Psych.* **19**, 61–73.

Moll, R. P. (1964). Drive and maturation effects in the development of consummatory behavior. *Psychol Reports* **15**, 295–302.

Morrell, F. (1961). Electrophysiological contributions to the neural basis of learning. *Physiol. Rev.* **41**, 443–494.

Moruzzi, G. and Magoun, H. W. (1949). Brain stem reticular formation and activation of the EEG. *Electroencephal. Clin. Neurophysiol.* **1**, 455–473.

O'Keefe, J. O. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.

Olds, J. (1955). Physiological mechanisms of reward. *In: Nebraska Symposium on Motivation* (Jones, M. R., Ed.). University of Nebraska Press, Lincoln.

Olds, J. (1958). Effects of hunger and sex hormone on self-stimulation of the brain. *J. Comp. Physiol. Psych.* **51**, 320–324.

Olds, J. (1977). *Drives and Reinforcements: Behavioral Studies of Hypothalamic Functions*. Raven Press, New York.

Pavlik, W. B. and Reynolds, W. F. (1963). Effects of deprivation schedule and reward magnitude on acquisition and extinction performance. *J. Comp. Physiol. Psych.* **56**, 452–455.

Pearce, J. M. and Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* **87**, 532–552.

Peters, A. Proskauer, C. C., Feldman, M. L. and Kimerer, L. (1979). The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. V. Degenerating axon terminals synapsing with Golgi impregnated neurons. *J. Neurocytol.* **8**, 331–357.

Reiss, D. (1970). Sidman avoidance in rats as a function of shock intensity and duration. *J. Comp. Physiol. Psych.* **73**, 481–485.

Rescorla, R. A. (1971). Variations in the effectiveness of reinforcement and nonreinforcement following prior inhibitory conditioning. *Learning and Motivation* **2**, 113–123.

Rescorla, R. A. and Wagner, A. R. (1972). A theory of Pavlovian conditioning. Variations in the effectiveness of reinforcement and nonreinforcement. *In: Classical Conditioning II: Current Research and Theory* (Black, A. H. and Prokasy, W. F., Eds.). Appleton–Century–Crofts, New York.

Reynierse, J. H. and Rizley, R. C. (1970). Relaxation and fear as determinants of maintained avoidance in rats. *J. Comp. Physiol. Psych.* **72**, 223–232.

Reynolds, W. F. and Pavlik, W. B. (1960). Running speed as a function of deprivation period and reward magnitude. *J. Comp. Physiol. Psych.* **53**, 615–618.

Robson, J. G. (1975). Receptive fields: Neural representation of the spatial and intensive attributes of the visual image. *In: Handbook of Perception*, Vol. 5 (Cartarette, E. C. and Friedman, M. P., Eds.). Academic Press, New York.

Rusinov, V. S. (1953). An electrophysiological analysis of the connecting function in the cerebral cortex in the presence of a dominant area. *Communications at the XIX International Physiological Congress, Montreal.*

Scott, J. W. and Pfaffmann, C. (1967). Olfactory input to the hypothalamus: Electrophysiological evidence. *Science* **158**, 1592–1594.

Seligman, M. E. P. and Hager, J. L. (Eds.) (1972). *Biological Boundaries of Learning*. Appleton–Century–Crofts, New York.

Seward, J. P., Shea, R. A. and Davenport, R. H. (1960). Further evidence for the interaction of drive and reward. *Am. J. Psych.* **73**, 370–379.

Seward, J. P., Shea, R. A. and Elkind, D. (1958). Evidence for the interaction of drive and reward. *Am. J. Psych.* **71**, 404–407.

Seward, J. P. and Proctor, D. M. (1960). Performance as a function of drive, reward, and habit strength. *Am. J. Psych.* **73**, 448–453.

Sheffield, F. D. (1951). The contiguity principle in learning theory. *Psychol. Rev.* **58**, 362–367.

Shepherd, G. M. (1974). *The Synaptic Organization of the Brain*. Oxford University Press, New York.

Sholl, D. (1956). *The Organization of the Cerebral Cortex*. Methuen, London.

Siegel, G. J., Albers, R. W., Katzman, R., Agranoff, B. W. (1976). *Basic Neurochemistry*. Second edition. Little, Brown, and Co., Boston.

Soechting, J. F. and Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *J. Neurosci.* **1**, 710–720.

Spence, K. W. (1956). *Behavior Theory and Conditioning*. Yale University Press, New Haven.

Squires, N. K., Donchin, E., Squires, K. C. and Grossberg, S. (1977). Redundant information in auditory and visual modalities: Inferring decision-related processes from the P300 component. *J. Exp. Psych.* **3**, 299–315.

Stein, L. (1958). Secondary reinforcement established with subcortical stimulation. *Science* **127**, 466–467.

Steinberg, M. S. (1970). Does differential adhesion govern self-assembly processes in histogenesis? Equilibrium configurations and the emergence of a hierarchy among populations of embryonic cells. *J. Exp. Zool.* **173**, 395–434.

Stolurow, L. M. (1951). Rodent behavior in the presence of barriers: II. The metabolic maintenance method; a technique for caloric drive control and manipulation. *J. Genetic Psych.* **79** 289–355.

Strouthes, A. (1965). Effect of CS-onset UCS-termination delay, UCS duration, CS-onset UCS-onset interval, and number of CS–UCS pairings on conditioned fear response. *J. Exp. Psych.* **69**, 287–291.

Swaab, D. F. and Schadé, J. P. (Eds.) (1974). *Integrative Hypothalamic Activity. Progress in Brain Research*, Vol. 41. Elsevier, Amsterdam.

Tapp, J. T. (1969). *Reinforcement and Behavior*. Academic Press, New York.

Thompson, R. F. (1967). *Foundations of Physiological Psychology*. Harper and Row, New York.

Tsumoto, T., Creutzfeldt, O. D. and Legendy, C. R. (1978). Functional organization of the corticofugal system from visual cortex to lateral geniculate nucleus in the cat. *Exp. Brain Res.* **32**, 345–364.

Valenstein, E. S. (1969). Biology of drives. *In: Neurosciences Research Symposium Summaries*, Vol. 3 (Schmitt, F. O., Melnechuk, T., Quarton, G. C. and Adelman, G., Eds.). M.I.T. Press, Cambridge, Mass.

Valenstein, E. S., Cox, V. C. and Kakolewski, J. W. (1970). Reexamination of the role of the hypothalamus in motivation. *Psychol. Rev.* **77**, 16–31.

Wagner, A. R. (1978). Expectancies and the priming of STM. *In: Cognitive Processes in Animal Behavior* (Hulse, S. H., Fowler, H. and Honig, W. K., Eds.). Erlbaum, Hillsdale, N.J.

Walter, W. G. (1964). Slow potential waves in the human brain associated with expectancy, attention, and decision. *Arch. Psychiat. Nervenkr.* **206**, 309–322.

Woolsey, C. N. (1958). Organization of somatic sensory and motor areas of the cerebral cortex. *In: Biological and Biochemical Bases of Behavior* (Harlow, H. F. and Woolsey, C. N., Eds.). University of Wisconsin Press, Madison.

Zaretsky, H. H. (1966). Learning and performance in the runway as a function of the shift in drive and incentive. *J. Comp. Physiol. Psych.* **62**, 218–221.

Zeaman, D. (1949). Response latency as a function of the amount of reinforcement. *J. Exp. Psych.* **39**, 446–483.

Zeki, S. M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J. Physiol.* **242**, 827–841.