Psychological Review
1982, Vol. 89, No. 5, 529-572

0033-295X/82/8905-0529$00.75

Processing of Expected and Unexpected Events During Conditioning and Attention: A Psychophysiological Theory

Stephen Grossberg
Department of Mathematics, Boston University

Some recent formal models of Pavlovian and instrumental conditioning contain internal paradoxes that restrict their predictive power. These paradoxes can be traced to an inadequate formulation of how mechanisms of short-term memory and long-term memory work together to control the shifting balance between the processing of expected and unexpected events. Once this formulation is strengthened, a unified processing framework is suggested wherein attentional and orienting subsystems coexist in a complementary relationship that controls the adaptive self-organization of internal representations in response to expected and unexpected events. In this framework, conditioning and attentional constructs can be more directly validated by interdisciplinary paradigms in which seemingly disparate phenomena can be shown to share similar physiological and pharmacological mechanisms. A model of cholinergic–catecholaminergic interactions suggests how drive, reinforcer, and arousal inputs regulate motivational baseline, hysteresis, and rebound, with the hippocampus as a final common path. Extinction, conditioned emotional responses, conditioned avoidance responses, secondary conditioning, and inverted U effects also occur. A similar design in sensory and cognitive representations suggests how short-term memory reset and attentional resonance occur and are related to evoked potentials such as N200, P300, and contingent negative variation (CNV). Competitive feedback properties such as pattern matching, contrast enhancement, and normalization of short-term memory patterns make possible the hypothesis testing procedures that search for and define new internal representations in response to unexpected events. Long-term memory traces regulate adaptive filtering, expectancy learning, conditioned reinforcer learning, incentive motivational learning, and habit learning. When these mechanisms act together, conditioning phenomena such as overshadowing, unblocking, latent inhibition, overexpectation, and behavioral contrast emerge.

Internal Problems of Some Conditioning Models

1. Merging Parallel Streams of Theory on Conditioning and Attention

This article compares and contrasts two parallel streams of theoretical progress in the conditioning and attention literature since 1968, using the article of Pearce and Hall (1980) as a basis for discussion. One stream was energized by such seminal articles as those of Estes (1969), Kamin (1968, 1969), Mackintosh (1971), Rescorla and Wagner (1972), and Wagner and Rescorla (1972). The great heuristic value of these articles stimulated new developments in such articles as those of Dickinson, Hall, and Mackintosh (1976), Frey and Sears (1978), Hall and Pearce (1979), Mackintosh (1976), Mackintosh, Bygrave, and Picton (1977), Mackintosh and Reese (1979), Sutton and Barto (1981), and Wagner (1976, 1978). The other stream is found in a series of my own articles (Grossberg, 1968, 1969a, 1969c, 1971, 1972a, 1972b, 1974, 1975, 1976a, 1976b, 1978a, 1978b, 1980, 1981b, 1982a, 1982b).

This is a good time to make this comparison because ideas from the two streams have gradually converged over the years. Once their remaining differences are resolved, both streams may be merged into a theoretical framework wherein conditioning, cognitive,

The preparation of this article was supported in part by the National Science Foundation (NSF IST-80-00257) and the Air Force Office of Scientific Research (AFOSR 82-0148).

Requests for reprints should be sent to Stephen Grossberg, Department of Mathematics, Boston University, 264 Bay State Road, Boston, Massachusetts 02215.

motivational, psychophysiological, and pharmacological data can be discussed in a unified fashion. In this framework, theoretical alternatives and predictions can be studied using interdisciplinary paradigms that can probe interactions that are opaque to more conventional experiments. Some experiments of this type will be summarized below.

2. The Processing of Expected and Unexpected Events in Short-Term Memory and Long-Term Memory

I suggest that various difficulties faced by the first stream are due to the fact that it does not adequately probe how mechanisms of short-term memory (STM) and long-term memory (LTM) influence the shifting balance between the processing of expected and unexpected events. These difficulties take several related forms: (a) Internal paradoxes exist within the theories. (b) No one theory can explain all the relevant data. In fact, no one theory can explain all the data explicable by any of the other theories. (c) The theories provide formal, as opposed to physical, models of the data. These models have no verifiable properties outside of the conditioning experiments they are constructed to explain. When some of these formal properties are interpreted as physical mechanisms, they are found to be either paradoxical or to have no external experimental support.

To clarify these assertions I will review concepts from Pearce and Hall (1980) as a source for this first stream. To illustrate how my approach overcomes these difficulties, I will review concepts from Grossberg (1980) as a source for the second stream, although the main concepts and mechanisms that I will need appeared in Grossberg (1971, 1972a, 1972b, 1975).

3. Some Internal Paradoxes

First I will review an internal paradox that lies at the heart of the Pearce and Hall (1980) theory. Pearce and Hall assert that "stimuli that fully predict their consequences will be denied access to the processor. . . . A stimulus is likely to be processed to the extent that it is not an accurate predictor of its consequences" (p. 538). Or, "a stimulus will gain access to the processor only when it . . . has been followed by a surprising event" (p. 540). One consequence of this position is that an unconditioned stimulus (US) that is an excellent predictor of food will not be processed even if no conditioned stimulus (CS) is present. Despite this implication, Pearce and Hall state that "stimuli such as the USs used in typical conditioning procedures are always likely to gain access to the processor" (p. 538). Pearce and Hall need this assumption because "conjoint processing of the CS and US representations results . . . in an increase in the ability of the CS to excite what we may call a 'US memory' " (p 542). This is the main conditioning event of their theory.

In an effort to embed their hypothetical processor into a broader theoretical perspective, Pearce and Hall analogize the processor to the limited-capacity STM system of human information-processing models. Given this processing interpretation, the Pearce and Hall model simultaneously implies that an expected US will not be stored in STM because it is expected and will be stored in STM because it is a US. One might try to escape this contradiction by claiming that the processor somehow knows the difference between a CS and a US. Even if one could overcome the problem of showing how the processor knows this difference, one would then be faced by the harder problem of showing how the processor changes its mind about a cue when the cue switches from CS to US status as a result of prior conditioning, and is thereupon used as a US in a secondary conditioning paradigm.

One can summarize this internal contradiction within the Pearce and Hall theory by saying that these authors have emphasized the processing of events that have unexpected consequences at the cost of implying paradoxes about the processing of events that have expected consequences.

Mackintosh (1975) developed a theory that emphasizes the processing of events that have expected consequences at the cost of falling into difficulties when explaining the processing of unexpected events. Pearce and Hall (1980) summarized Mackintosh's position as follows: "He suggested that the associability (α) of a stimulus will increase if it predicts reinforcement more accurately

than other stimuli present in the situation but will decrease if it predicts reinforcement less accurately" (p. 536). This hypothesis was made to explain how conditioning of a Cue X is blocked if X is presented on compound trials AX after prior conditioning of Cue A has occurred. Mackintosh's hypothesis explains blocking by claiming that X does not condition well because A is a better predictor of the US due to its prior conditioning trials.

Mackintosh's hypothesis is, however, incompatible with the fact that the Cue X conditions normally on the first compound trial (Kamin, 1968; Mackintosh et al., 1977; Rescorla & Wagner, 1972). This experimental result contradicts the hypothesis because on the first compound trial, Cue X is a worse predictor of the US than it is on the second compound trial. Why does the Cue X condition normally on the first trial if on this trial it is the worst possible predictor of the US, having never before been correlated with the US?

To escape this contradiction Mackintosh simply assumes that only the intensity of X, not its predictability, influences its associability on the first compound trial. Whether one considers this an internal paradox of his theory or an ad hoc restatement of the data is a matter of taste.

Despite this difficulty, Pearce and Hall (1980) write that "the success of Mackintosh's model . . . convinces us that the principle it embodies—the modification of CS associability as a result of the consequences of one trial influencing conditioning on the next—must be a part of any successful theory" (p. 537).

Wagner (1976, 1978) has attempted to give Mackintosh's ad hoc assumption a physical basis by assuming "that the associability of a CS is inversely related to the strength of an association between the CS and the context" (Pearce & Hall, 1980, p. 549). Pearce and Hall (1980) criticize Wagner's concept by noting that "it is very difficult to see how a surprising shock omission . . . or shock increase . . . after a CS can reduce the strength of the association between the CS and the context and thus restore associability" (p. 549). Their criticism does not distinguish between the processing of an event that is unexpected within a given context and the processing of an event that is followed by an unexpected US. I will suggest below how such distinctions tend to be blurred within these formal models and how they can be clarified using a physically based model.

4. The Need for Behaviorally Unobservable Mechanisms

The seriousness of the dilemma into which conditioning data have driven the formal models can be appreciated from the following considerations: Mackintosh (1975) says that events that have expected consequences are processed, whereas Pearce and Hall (1980) say that events that have unexpected consequences are processed. Both viewpoints are, moreover, supported by unimpeachable data. If the data support the idea that both expected and unexpected events are processed, then why have the formal theories avoided this conclusion?

The statement that both expected and unexpected events are processed can easily become predictively vacuous in a formal model, because such a model cannot easily distinguish the sense in which expected and unexpected events are processed in different ways. The alternative conclusion—that expected and unexpected events are processed in the same way—is dangerously close to saying that all events are processed in the same way, which is patently false.

Conditioning data need a theory that can avoid these fatal pitfalls. Such a theory must explain the sense in which expected and unexpected events are processed by different mechanisms. It must carefully delineate the properties of these mechanisms to avoid becoming vacuous or patently false. It must show how these properties can be empirically tested. Because expected and unexpected events are all just events on the behaviorally observable level, such a theory needs to establish a link with the behaviorally unobservable structures within which these processing distinctions can be physically interpreted and validated.

5. Causality Violation on the Behaviorally Observable Level

The step toward theories that invoke behaviorally unobservable processes in a sub-

stantive way runs against the grain for many psychologists today. This is true despite the fact that as a body of psychological data becomes more mature and quantitative, it bears more sharply on the behaviorally unobservable mechanisms that generate these data.

An important experiment that further demonstrates that conditioning data have reached this level of maturity was conducted by Mackintosh et al. (1977) and is reviewed by Pearce and Hall (1980). I will indicate below that an interpretation of these data using only behaviorally observable variables violates the causality of the conditioning process. To reject causality is tantamount to denying the very existence of a predictive conditioning theory. For my purposes, I will summarize only one aspect of this experiment.

In Part 1 of the experiment, all rats experienced four trials on which a light (CS) was followed by a shock (US). In Part 2 of the experiment, two groups of rats received an additional single compound light–tone trial. In one group (Group 1) the light–tone compound was followed by a single shock. In the other group (Group 2), the light–tone compound was followed by two successive shocks that were presented 10 sec apart. A recall trial with the tone alone showed essentially identical fear conditioning to the tone in both groups. In other words, the second shock seems not to have affected tone conditioning.

The remarkable feature of the experiment becomes apparent when one considers two other groups of rats tested in Part 3 of the experiment. One of these groups (Group 3) received the same training as Group 1 did plus an additional compound light–tone trial followed by a single shock before a recall trial with the tone CS. The other group (Group 4) received the same training as did Group 2 plus an additional compound light–tone trial followed by a single shock before a recall trial with the tone CS. In other words, Part 3 of the experiment simply added an identical learning manipulation onto the Group 1 and 2 learning paradigms, which by themselves elicited the same reaction to the tone. Remarkably, the tone exhibited better fear conditioning for Group 4 than for Group 3.

This is a fascinating experimental finding. How can a test after identical second compound trials have different effects if tests after different first compound trials had identical effects? This experiment seems to violate causality on the behaviorally observable level and forces us to turn to behaviorally unobservable mechanisms for an explanation.

6. Some Unpredicted Data

The internal contradictions within the formal models are associated with predictive limitations. Pearce and Hall (1980) note that they cannot explain the following phenomena: (a) If two CSs differ markedly in their intensity or salience, the more salient cue can overshadow the less salient cue, but not conversely (p. 541; Mackintosh, 1976). (b) Overshadowing effects can sometimes occur when only a single trial of compound conditioning is given (p. 541; Mackintosh, 1971; Mackintosh & Reese, 1979). (c) The associability of a stimulus followed by surprising food remains high for future conditioning involving food but not for future conditioning involving shock (p. 550; Dickinson & Mackintosh, 1979). (d) Low associability of a stimulus can be restored by presenting it in a novel context (p. 550; Dexter & Merrill, 1969; Lantz, 1973; Lubow, Rifkin, & Alek, 1976). (e) The occurrence of a surprising event soon after a conditioning trial can influence learning on that trial (p. 550; Kremer, 1979; Wagner, Rudy, & Whitlow, 1973). All of these phenomena can be explained by my theory. See Sections 33–45.

7. Formal Versus Physical Concepts: A Second Type of LTM?

Formal models have an advantage over heuristic data analysis in that they commit one's thinking to a more precise, and therefore disconfirmable, set of concepts. Formal models also need to be compared with other related concepts in the theoretical literature to check their tenability as physical constructs. The main conditioning equations of the Pearce and Hall (1980) theory lead to a major qualitative conclusion when this comparison is made, namely, that "associability" is controlled by a form of LTM distinct from the LTM that is encoded by associative strength.

This conclusion follows from the Pearce and Hall associative equations:

$$\Delta V_A^{(n)} = S_A^{(n)} \alpha_A^{(n)} \lambda^{(n)} \quad (1)$$

and

$$\alpha_A^{(n)} = |\lambda^{(n-1)} - V_T^{(n-1)}| \quad (2)$$

or Equation 1 and

$$\alpha_A^{(n)} = \frac{1}{c} \sum_{k=n-c}^{n-1} |\lambda^{(k)} - V_T^{(k)}|. \quad (3)$$

Equation 1 says that the change $\Delta V_A^{(n)}$ in the associative strength $V_A^{(n)}$ of Event A on trial n depends on the product of CS intensity $S_A^{(n)}$, associability $\alpha_A^{(n)}$, and US intensity $\lambda^{(n)}$ on trial n. Equation 2 says that the associability on trial n depends on—and *a fortiori* remembers—the difference between the US intensity and the aggregate associative strength $V_T^{(n-1)}$ on trial $n-1$. Equation 3 extends this notion to the idea that associability has a memory that extends c (>1) trials into the past.

All of the formal theories make some assumption concerning how events on a given trial influence associability on later trials. If these formal theories have physical validity, we must therefore determine whether two distinct types of LTM exist. If so, how does associability physically compute the difference between the US intensity and the aggregate associative strength of all experimental events despite the fact that these quantities are distributed across time and representational space? What sort of intervening processes can gather, add, or subtract all these data at each CS representation? Is there any physical evidence that such processes exist? If these formulations have no physical support as they stand, is it because the formal theories have lumped too many processes into a single equation?

8. Overexpectation, STM Signaling, LTM Gating, and Drive Representations

The scope of this article does not permit a derivation of my theory from first principles. I can, nonetheless, indicate how a bridge can be constructed from the Rescorla–Wagner framework, on which Pearce and Hall have built, toward my own framework. I do this by analyzing an internal contradiction in the way Wagner and Rescorla (1972) use their associative equation to explain the phenomenon of "overexpectation." This contradiction arises because Wagner and Rescorla's intuitive explanation is incompatible with their formal associative equation. By explicating their intuitive explanation within their own formalism, I am led to a distinction between STM and LTM effects and to the concept of a drive representation.

The phenomenon of overexpectation was observed in an experiment wherein four groups of rats received cue combinations in a fear conditioning (CER) paradigm. On each conditioning trial a 2-min flashing light (A) and/or a 2-min tone (B) preceded a .5-sec .5-ma shock US. One group (A+B+) received 40 reinforced presentations of each cue, A and B, presented separately. Group A+B+/AB+ received, in addition, 40 more reinforced presentations of the AB compound. Two comparison groups received either 80 or 120 reinforced trials with only the compound stimulus. All groups then received 16 nonreinforced presentations of the AB compound while barpressing for food to test the degree of response suppression produced by the compound. Group A+B+ showed more conditioned suppression of bar pressing for food during unreinforced test trials than any of the other groups.

Wagner and Rescorla (1972) discuss these results in terms of their associative equations

$$\Delta V_A = \alpha(\lambda - V_\Sigma) \quad (4)$$

and

$$\Delta V_B = \alpha(\lambda - V_\Sigma), \quad (5)$$

where V_A and V_B are the associative strengths of A and B, respectively, $V_\Sigma = V_A + V_B$, and λ is the US strength. They claim that when only Cue A is followed by shock, V_A will approach λ because then $V_\Sigma = V_A$. Similarly, when only Cue B is shocked, V_B will approach λ because then $V_\Sigma = V_B$. If, as in the experiment A+B+/AB+, the compound AB is thereupon shocked, then $V_\Sigma = V_A + V_B$. Consequently, both V_A and V_B will decrease when AB training begins because

$$\lambda - V_\Sigma \cong \lambda - 2\lambda < 0 \quad (6)$$

at the onset of compound training.

This explanation is inconsistent with

Equations 4 and 5 for the following reason: On the A+B+ trials, V_A and V_B are each assumed to approach λ even before compound AB trials begin. Because, by definition, $V_\Sigma = V_A + V_B$, this cannot happen because as soon as $V_\Sigma \geq \lambda$, $\Delta V_A \leq 0$ and $\Delta V_B \leq 0$. In other words, the Wagner and Rescorla (1972) model, when consistently applied, is inconsistent with the fact that more suppression can occur after A+B+ trials than after AB+ trials.

Rescorla and Wagner (1972) clearly intended that V_B should be irrelevant when only Cue A is presented and that V_A should be irrelevant when only Cue B is presented. Such a concept is needed to eliminate the effect of V_B on the V_A asymptote during A trials, and of V_A on the V_B asymptote during B trials. The switching on and off of an event's momentary relevance occurs rapidly on a trial-by-trial basis, whereas the growth of the associative strengths V_A and V_B is slowly varying across trials. Let us call the distinction between the rapid modulation of an event's activity and the slow changes in its associative strengths the difference between STM and LTM.

Had Wagner and Rescorla (1972) explicitly faced this processing implication of their own data, they might have redefined V_Σ as

$$V_\Sigma = S_A V_A + S_B V_B , \qquad (7)$$

where S_A and S_B are the signals elicited by the active STM representations of Cues A and B, respectively. Then, If only A is presented, $S_A = 0$ and V_Σ depends only on V_B, whereas if only B is presented, $S_B = 0$ and V_Σ depends only on V_A.

Term $S_A V_A$ in Equation 7 can be physically interpreted as follows: The STM-activated signal S_A reads out the LTM trace V_A via a *gating*, or multiplicative, action $S_A V_A$. Such an LTM gating action appears prominently in my work. It is, for example, crucial in my approach to serial and paired-associate verbal learning (Grossberg, 1969b; Grossberg & Pepe, 1971), free recall (Grossberg, 1978a, 1978b), and cognitive development (Grossberg, 1976a, 1976b, 1978b, 1980), as well as in my studies of conditioning and attention (Grossberg, 1968, 1969a, 1969b, 1969c, 1971, 1972a, 1972b, 1974, 1975, 1976a, 1976b, 1978a, 1978b, 1980, 1981b, 1982a, 1982b).

Once the modification in Equation 7 is accepted, it becomes clear that the model must undergo a more major revision. This is true because the left-hand side and the right-hand side of Equation 7 are not dimensionally the same. The right-hand side can fluctuate rapidly through time with S_A and S_B, whereas the left-hand side is a slowly varying associative strength. This observation could have already been made about the term $\lambda - V_\Sigma$ in Equations 4 and 5, because US intensity λ is a rapidly varying (STM) quantity, whereas V_Σ is a slowly varying (LTM) quantity. Replacing V_Σ by $S_A V_A + S_B V_B$ in $\lambda - V_\Sigma$ avoids the problem of mixing apples with oranges if we interpret $\lambda - S_A V_A - S_B V_B$ as the amount by which a commonly shared STM representation is activated by the combined effects of λ and the LTM-gated signals $S_A V_A$ and $S_B V_B$.

This commonly shared STM representation cannot be the separate representations of either Cue A or Cue B. Moreover, activation of this new representation depends on the choice of reinforcer, because an associative strength learned with respect to a shock is not the same as an associative strength learned with respect to a food reinforcer. Even within the Rescorla–Wagner model, V_Σ feels the influence of a particular US's intensity λ. In my theory this new type of STM representation is called a *drive representation.*

An internal analysis of the Rescorla–Wagner equation has hereby distinguished sensory representations for cues such as A and B from drive representations corresponding to distinct reinforcing actions such as food, fear, and sex. Because each Cue A might be conditioned to any of several drive representations, we need to study how the *pattern* of LTM associative strengths V_{ij}, leading from the ith sensory representation to the jth drive representation, evolves through time. Once we accept the fact that the ith sensory representation can read out a pattern of LTM-gated signals $S_i V_{ij}$ across several drive representations (indexed by j), we need to discover an STM decision rule whereby incompatible drive representations can generate consistent observable behavior. We also need to discover a law for the selective change of V_{ij} due to the STM signal S_i of the ith CS and the intensity λ_j of the jth US.

9. Secondary Conditioning Implies That Cue and Drive Representations Are Distinct

Now that we have in mind sets not only of CSs but also of USs, we can use the fact that prior conditioning can transform a CS into the US of a later secondary conditioning experiment to constrain this law. In particular, the asymmetric role of US intensity λ_j and of CS intensity S_i in the modified Rescorla–Wagner equation,

$$\Delta V_{ij} = \alpha_{ij}(\lambda_j - \sum_k S_k V_{kj}), \qquad (8)$$

shows that this equation cannot be strictly correct. I suggest that this problem of the Rescorla–Wagner framework is the reason why Pearce and Hall need to assume that two types of LTM exist (Section 33). Equation 8 also includes the Widrow-Hoff equation on which Sutton and Barto (1981) build.

The same argument shows that a secondary US representation is not a drive representation, because a CS representation is not a drive representation. Hence, conditioning from a CS to a drive representation is not the same process as conditioning from a CS to a US representation. This conclusion runs counter to the Pearce and Hall (1980) assertion "that the amount of learning is determined by the amount of simultaneous processing that the representations of the CS and US receive in the processor" (p. 550).

In light of the above argument, it is not clear what "the processor" might physically represent, because CS representations and drive representations are qualitatively distinct concepts. In fact, Pearce and Hall (1980) note their model's inability to explain the Dickinson and Mackintosh (1979) data on selective effects of distinct reinforcers on associability and go on to say, "One way for our model to accommodate this result is to propose that there are separate processors for learning about different reinforcers such as food and shock" (p. 550). A large body of data other than that of Dickinson and Mackintosh (1979) also suggests such a concept. My articles (Grossberg, 1971, 1972a, 1972b, 1975, 1982a) review some of these data in light of the drive representation concept.

This type of internal analysis of the Rescorla–Wagner framework can be continued, but the breakdown of Equation 8 indicates that some new theoretical principles are needed to go much further. The above theoretical exercise nonetheless clarifies my contention that the demand for explicit processing descriptions—within any theoretical framework—rapidly leads either to important new concepts or to unforeseen contradictions.

Some General Psychophysiological Concepts

10. An Alternative Processing Framework: Complementary Attentional and Orienting Subsystems

Having summarized some difficulties of one theoretical stream, I will compare the two streams—notably their explanations of expectancies, extinction, and STM priming—after my review of the second stream is complete. To start the review, I will sketch in broad strokes the general framework of my theory. Then I will review in more precise terms the several design principles and mechanisms that I need to quantify this framework.

In my theory an interaction between two functionally complementary subsystems is needed to process expected and unexpected events (Grossberg, 1975). A precursor of this concept is developed in the distinguished psychophysiological article of Routtenberg (1968) on the "two-arousal hypothesis." My conception of these two subsystems will be seen to deviate from Routtenberg's view in several basic ways (Section 27).

Expected events are processed within a consummatory, or attentional, subsystem. This subsystem establishes ever more precise internal representations of and responses to expected cues. It also builds up the learned expectations that are used to characterize operationally the sense in which expected cues are expected. The attentional subsystem is, however, incapable of adapting to unexpected environmental changes. Left to its own devices, it would elicit ever more rigid, even perseverative, reactions to the environment, much as hippocampectomized rats do not orient to a novel stimulus while they are indulging in consummatory activity, such as running toward a reward. Such rats cannot "shift attention during the presentation of a novel stimulus or in a mismatch situation"

(O'Keefe & Nadel, 1978, p. 250). The second subsystem is an orienting subsystem that overcomes the rigidity of the attentional subsystem when unexpected events occur and enables the attentional subsystem to adapt to new reinforcement and expectational contingencies.

Part of the difficulty in understanding conditioning and attentional data is due to the fact that these two subsystems interact in a subtle fashion. I will review in the following sections how both expected and unexpected events start to be processed by the attentional subsystem. When an unexpected event mismatches an active expectancy within this subsystem, the orienting subsystem is disinhibited. The orienting subsystem acts to rapidly reset STM within the attentional subsystem as it simultaneously energizes an orienting response.

By contrast, an expected event matches an active expectancy within the attentional subsystem. This matching process amplifies the STM activity patterns that are currently active within the attentional subsystem. These amplified, or resonant, STM activities inhibit the orienting subsystem as they simultaneously drive adaptive LTM changes, including the learning of new expectancies, internal representations (chunks), and habits.

11. The Stability–Plasticity Dilemma and Evoked Potential Correlates

The complementary attentional and orienting subsystems, indeed all the mechanisms that I will use, arise as the solution to a fundamental design problem concerning the self-organization (e.g., development, learning) of new internal representations (Grossberg, 1976a, 1976b, 1978b, 1980, 1982a, 1982b). I call this problem the stability–plasticity dilemma.

The stability–plasticity dilemma concerns how internal representations can maintain themselves in a stable fashion against the erosive effects of behaviorally irrelevant environmental fluctuations yet can nonetheless adapt rapidly in response to environmental fluctuations that are crucial to survival. How does a network as a whole know the difference between behaviorally irrelevant and relevant events even though its individual cells, or nodes, do not possess this knowledge? How does a network transmute this knowledge into the difference between slow and fast rates of adaptation, respectively? Classical examples of the stability–plasticity balance are found in the work of Held and his colleagues on rapid visual adaptation to discordant visuomotor data in adults (Held, 1961, 1967; Held & Hein, 1963) and in the work of Wallach and his colleagues on rapid visual adaptation to discordant cues for the kinetic depth effect and cues for retinal disparity (Wallach & Karsh, 1963a, 1963b; Wallach, Moore, & Davidson, 1963).

Because of the fundamental nature of the stability–plasticity dilemma, the mechanisms from which the two complementary subsystems are built have properties that imply psychophysiological, neurophysiological, and pharmacological predictions. For example, on the psychophysiological level, the disinhibition of the orienting subsystem due to an expectancy mismatch is suggested to correspond to the mismatch-negativity component of the N200 evoked potential complex. The STM reset in the attentional subsystem is suggested to correspond to a P300 evoked potential. The origin of the mismatch-negativity component in the orienting subsystem and its role in generating a P300 suggests a relationship between the P300 and the orienting reaction. The resonant STM activity that derives from an expectancy match in the attentional subsystem is suggested to correspond to the processing-negativity component of the N200 evoked potential complex.

This psychophysiological interpretation leads to a number of interdisciplinary predictions (Grossberg, 1982a). For example, in Section 35, I suggest that the tone on the second compound trial in Group 4 of the Mackintosh, Bygrave and Picton (1977) experiment is more unexpected than the tone on the second compound trial in Group 3, and should therefore elicit a larger P300 evoked potential.

Sections 12–25 review the concepts that I use to mechanize the attentional and orienting subsystems. Then Sections 26–51 use these concepts to explain conditioning and attentional data and to compare my theory with the formal models.

12. Gated Dipoles

The gated dipole design is needed to reset STM. In the present theory the term *STM* refers collectively to the suprathreshold activities of STM traces. An STM trace is computed at a network node where it equals the average potential of the cell, or cell population, that is represented by the node. An STM trace can passively decay at a node, but the important operations in the theory transform these traces in ways other than by passive decay. In particular, STM reset refers to a rapid change in STM, notably the rapid shutting off of activity at a subset of previously active nodes and the rapid turning on of activity at a subset of previously inactive nodes. The STM activities at different types of nodes represent different psychological processes. The gated dipoles to be discussed below are, for example, assumed to occur both in cognitive and in motivational networks. Activity of a gated dipole node in a cognitive network may represent the occurrence of a certain perceptual feature or a certain temporal ordering of events in an experiment. Activity of a gated dipole node in a motivational network may, by contrast, measure the level of perceived fear or relief. The theory suggests that similar formal properties obtain in gated dipoles wherever these dipoles are placed in a network. The theory also suggests how to place these mechanisms in different types of networks and how to interpret their formal properties in these different contexts.

The gated dipole design shows how slowly accumulating transmitter substances gate the signals in parallel network pathways before these pathways compete to elicit net on-cell or off-cell STM responses (Figure 1). These responses include a sustained on-response to cue onset and a transient antagonistic off-response, or rebound, to either cue offset or to arousal onset. The off-reactions drive the STM reset.

One way to motivate the antagonistic rebound concept is to ask, How does offset of an event act as a cue for a learned response? For example, suppose that I wish to press a lever in response to the offset of a light. If light offset simply turned off the cells that code for light being on, then there would be no cells whose activity could selectively elicit the lever-press response after the light was turned off. Light offset must also selectively turn on cells that will transiently be active after the light is shut off. The activity of these off-cells (the cells that are turned on by light offset) can then activate the motor commands leading to the lever press. Let us call the transient activation of the off-cell by cue offset antagonistic rebound.

In a reinforcement context I claim that such an antagonistic rebound is the basis for a relief reaction (Denny, 1971) upon offset of a sustained fear-eliciting cue. In a perceptual context I claim that such an antagonistic rebound is the basis for a negative aftereffect upon offset of a sustained image (Brown, 1965, p. 483; Helmholtz, 1866, 1866/1962).

13. Antagonistic Rebound to Cue Offset

I will now describe a minimal model capable of eliciting a sustained on-response to onset of a cue and a transient antagonistic rebound to offset of the cue. The intuitive postulates that led to the model's original derivation are given in Grossberg (1972b). An alternative derivation is given in Grossberg (1980, Appendix E). An extended discussion and mathematical analysis of the gated dipole is found in Grossberg (1981b, 1982a). Herein I will merely provide an intuitive description of a gated dipole.

Consider Figure 1. In Figure 1 (a), a nonspecific arousal input I is delivered equally to both the on-channel and the off-channel, whereas a test input J (e.g., light or shock) is delivered only to the on-channel. These inputs activate the potentials X_1 and X_2, which create signals S_1 and S_2 in the on-channel and off-channel respectively. Because $I + J > I$, $X_1 > X_2$, and consequently, $S_1 > S_2$. What happens next is crucial.

The square synapses are assumed to contain chemical transmitters Z_1 and Z_2, respectively. Each transmitter slowly accumulates to a target level. The slow accumulation rate is essential to the model's properties. The target level is achieved by a constant transmitter production rate that is reduced by feedback inhibition proportional to the transmitter concentration. When a signal S_1 reaches the synaptic knobs containing

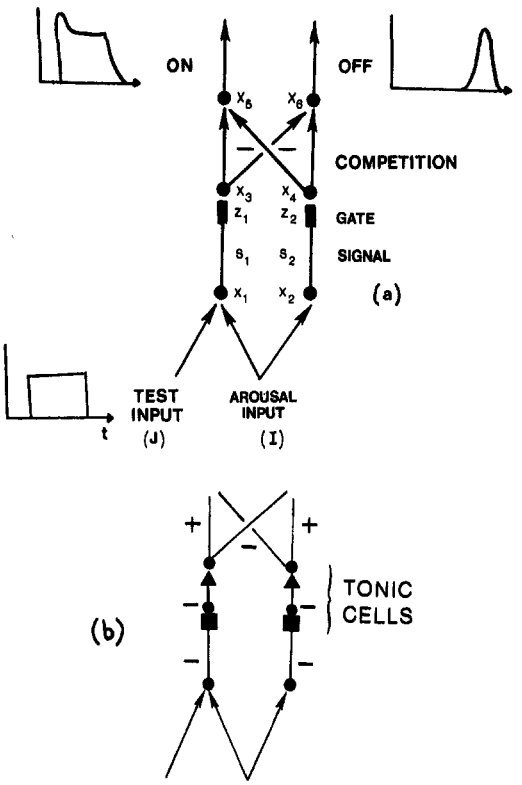

Figure 1. Two examples of gated dipoles. In Panel a the phasic input J and the arousal input I add in the on-channel, thereby activating the short-term memory trace X_1. The arousal input I also perturbs the short-term memory trace X_2 in the off-channel. Consequently, $X_1 > X_2$. Then X_1 and X_2 elicit signals $f(X_1)$ and $f(X_2)$ in their respective pathways. Because $X_1 > X_2, f(X_1) > f(X_2)$ also. 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 short-term memory traces X_5 and X_6 then satisfy $X_5 > X_6$. Each short-term memory trace elicits an excitatory signal down its own pathway, and an inhibitory signal to the other pathway. The net effect after competition takes place 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 Panel 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 that 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 thereby excite it.

Z_1, transmitter is released at a rate proportional to $T_1 = S_1Z_1$. The multiplicative effect of Z_1 on S_1 to yield T_1 is called transmitter gating of the signal S_1. The gating law just says that S_1 and Z_1 interact via mass action to elicit T_1. In particular, if either $S_1 = 0$ or $Z_1 = 0$, then $T_1 = 0$.

One proves that if $S_1 > S_2$, then $T_1 > T_2$. That is, transmitter is released at a faster rate by larger signals. Consequently, potential X_3 exceeds potential X_4. These potentials then emit competing signals. Potential X_5 wins the competition over X_6 and emits output signals that are the on-reaction of the network.

So far everything seems quite elementary. Only now do we exploit the slow accumulation rate and the transmitter gating law to show how a transient antagonistic rebound is generated by rapid offset of J.

The faster transmitter depletion rate in the on-channel than in the off-channel when J is on implies that $Z_1 < Z_2$, despite the fact that $S_1Z_1 > S_2Z_2$. When J is shut off, both channels receive the equal arousal input I. The potentials X_1 and X_2 rapidly equalize, as do the signals S_1 and S_2. By contrast, the inequality $Z_1 < Z_2$ persists because transmitter accumulates slowly. Thus, right after J shuts off, $S_1Z_1 < S_2Z_2, X_3 < X_4$, and the off-channel wins the signal competition. An antagonistic rebound is thereby initiated.

The rebound is transient because the transmitters gradually respond to the equal signals I by reaching a common level $Z_1 = Z_2$. Then $S_1Z_1 = S_2Z_2$, and the competition shuts off the rebound.

There exist many variations on the gated dipole theme. Figure 1 (b) points out that the slow transmitters can be inhibitory transmitters within a two-synapse disinhibitory pathway rather than excitatory transmitters within a one-synapse pathway. I interpret dopamine or noradrenaline to be the slow inhibitory transmitters in motivational and cognitive dipoles. The other inhibitory transmitter is often interpreted as gamma aminobutyric acid (Groves, Young, & Wilson, 1978). The disinhibitory concept rationalizes many effects of drugs such as amphetamine, chlorpromazine, 6-hydroxydopamine, and monoamine oxidase inhibitors on behavior (Grossberg, 1972b, 1982a). A single cell, rather than an intercellular network as in Figure 1, can also act like a gated dipole. Such a dipole is suggested to exist in vertebrate photoreceptors (Carpenter & Grossberg, 1981) wherein calcium is suggested to

act as the gating chemical. A full understanding of the gated dipole concept requires that we be able to distinguish the varied anatomical substrates of gated dipoles from their commonly shared functional properties.

14. Antagonistic Rebound to Arousal Onset

A surprising fact about gated dipoles is that a sudden arousal increment ΔI can trigger an antagonistic rebound despite the fact that both the on-channel and the off-channel receive equal arousal inputs (Figure 2). This mathematical property forced me to realize that an unexpected event, by triggering a burst of nonspecific arousal, could disconfirm an on-reaction by rapidly and selectively inhibiting it, thereby resetting STM.

I should be more precise about this property of arousal, because this precision has important implications for my explanation of overshadowing. The following remarkable property holds in the gated dipole of Figure 1 (a) if the signals S_1 and S_2 are linear functions of their respective inputs. The off-rebound size in response to a sustained input J and a sudden arousal increment of size ΔI above the previous arousal level I is

$$\text{Off} = \frac{ABJ(\Delta I - A)}{(A + I + J)(A + I)}, \tag{9}$$

where A and B are positive constants. Note that a positive off-reaction occurs only if $\Delta I > A$. This criterion is independent of J, which means that an arousal increment ΔI that is sufficiently large to rebound any dipole will be large enough to rebound all dipoles

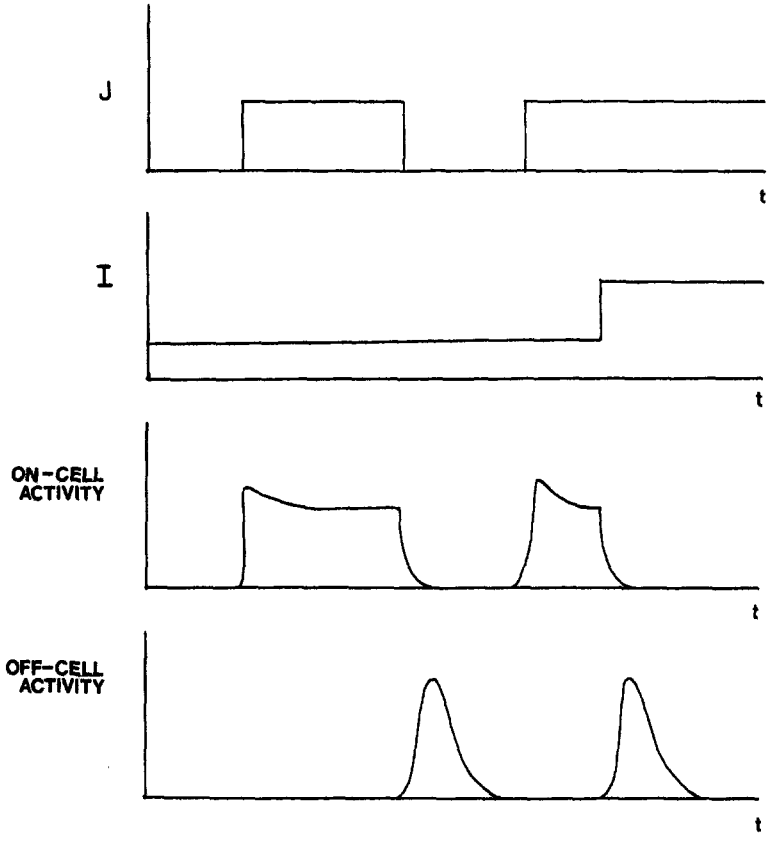

Figure 2. On and off responses of gated dipoles. After a gated dipole's transmitters equilibrate to an on-channel phasic input J, an antagonistic rebound or off-response can be generated by either rapidly shutting off the phasic input J or rapidly turning up the level of nonspecific arousal I. The latter type of rebound can reset short-term memory in response to an unexpected event that triggers a momentary arousal burst from the orienting subsystem.

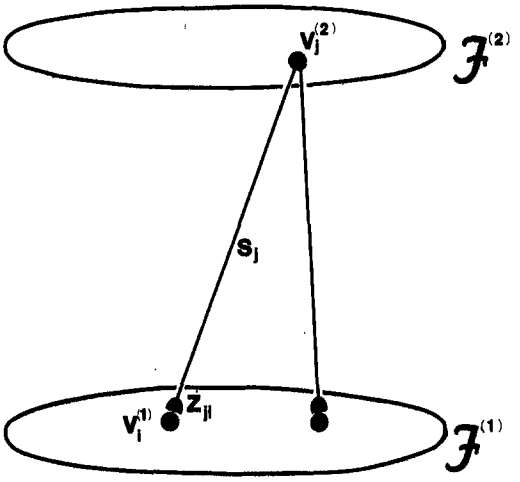

Figure 3. Short-term memory activity at a population $V_j^{(2)}$ in $F^{(2)}$ releases a signal S_j. This signal is gated by the long-term memory trace Z_{ji} on its way to a population $V_i^{(1)}$ in $F^{(1)}$. All these gated signals add at $V_i^{(1)}$ to generate a total input $E_i = \sum_j S_j Z_{ji}$. The pattern $E = (E_1, E_2, \ldots, E_m)$ of feedback signals is an expectation.

in a field. More precisely, if the arousal increment is large enough, then all active dipoles will simultaneously be rebounded. This is because the size of the rebound is an increasing function of the on-input J and equals 0 if $J = 0$. Thus, rebound size is selective despite the fact that all active dipoles can be rebounded at once. I identify activation of the arousal source which resets the STM of a dipole field with the mismatch-negativity component of the N200 evoked potential complex. The STM reset event itself is identified with the P300 evoked potential.

Given that a sudden burst of nonspecific arousal can selectively reset a field of on-cells and off-cells, we need to consider how this arousal level is regulated. The theory explicates the idea that surprising events are arousing.

15. What Is an Expectation?

To discuss a surprising or unexpected event, we need to define what an expectation is and how it is computed. In my theory several levels of network processing interact to build up new internal representations. To fix ideas, consider two successive levels $F^{(1)}$ and $F^{(2)}$ in this network hierarchy. Each active node in $F^{(2)}$ is capable of sending feedback signals to a subset of nodes across $F^{(1)}$. Before these signals reach their target cells, they are gated by LTM traces that exist at the ends of the signal pathways (Figure 3). This gating action multiplies the signal and the LTM trace, just like the gating action of the transmitters in a gated dipole. In more precise terms, denote the feedback signal from the jth node $V_j^{(2)}$ in $F^{(2)}$ to the ith node $V_i^{(1)}$ in $F^{(1)}$ by S_j. Denote the LTM trace in this feedback pathway by Z_{ji}. Then the net signal received by $V_i^{(1)}$ from $V_j^{(2)}$ is $S_j Z_{ji}$. All these signals add up to $V_i^{(1)}$ to generate a total signal

$$E_i = \sum_j S_j Z_{ji} \tag{10}$$

at $V_i^{(1)}$. The pattern

$$E = (E_1, E_2, \ldots, E_m) \tag{11}$$

of total feedback signals from $F^{(2)}$ to $F^{(1)}$ is identified with an expectation. An expectation is not an LTM trace or a family of LTM traces. It is not defined with respect to a single cell. It is a feedback pattern derived from LTM-gated signaling across the entire network. Because the signals S_j depend on the momentary STM activities of the nodes in $F^{(2)}$, the expectation can quickly change even if there are no changes in the LTM traces. The signal S_j is an STM probe that reads out the LTM pattern $(Z_{j1}, Z_{j2}, \ldots, Z_{jm})$ as part of the expectation E.

The pattern E is called an expectation because the LTM trace Z_{ji} can change when the signal S_j from $V_j^{(2)}$ and the STM activity $X_i^{(1)}$ of $V_i^{(1)}$ are active long enough for the slowly varying LTM trace to respond to them. When this happens, the pattern $(Z_{j1}, Z_{j2}, \ldots, Z_{jm})$ of LTM traces can learn the pattern $(X_1^{(1)}, X_2^{(1)}, \ldots, X_m^{(1)})$ of STM activities (Grossberg, 1967, 1969c, 1972c, 1974). As the STM pattern across $F^{(1)}$ is encoded by $V_j^{(2)}$'s LTM traces, it becomes the pattern that $V_j^{(2)}$ expects to find at $F^{(1)}$ when it is active. Later STM activation of $V_j^{(2)}$ reads this LTM pattern into the expectation E via the gating action of the LTM traces Z_{ji} on the signal S_j. This gated signal pattern equals E only when $V_j^{(2)}$ is the only active node in $F^{(2)}$. When more than one node is active across $F^{(2)}$, E is a weighted average of the

LTM patterns of all the active cells (Grossberg, 1968, 1976a, 1976b, 1980).

Neurophysiological evidence for the existence of such feedback expectancies, or templates, can be found in the distinguished work of Freeman (1975, 1980, 1981) on the olfactory system.

16. Unexpected Events Trigger a Mismatch-Modulated Arousal Burst

Having defined an expectation, I can now more easily describe how an unexpected event at $F^{(1)}$ triggers an arousal burst to $F^{(2)}$, which thereupon resets STM across $F^{(2)}$ via selective antagonistic rebounds.

- At a moment when the feedback expectation E is active across $F^{(1)}$ (Figure 4 [a]), suppose that an external event causes a feedforward input pattern U to be received by $F^{(1)}$ (Figure 4 [b]). Suppose that U mismatches E (in a sense that is defined in Section 24). In my theory such a mismatch rapidly inhibits STM activity across $F^{(1)}$ (Figure 4 [c]). Attenuating the STM activity in $F^{(1)}$ eliminates the inhibitory signal that $F^{(1)}$ delivers to the orienting, or arousal, sub-

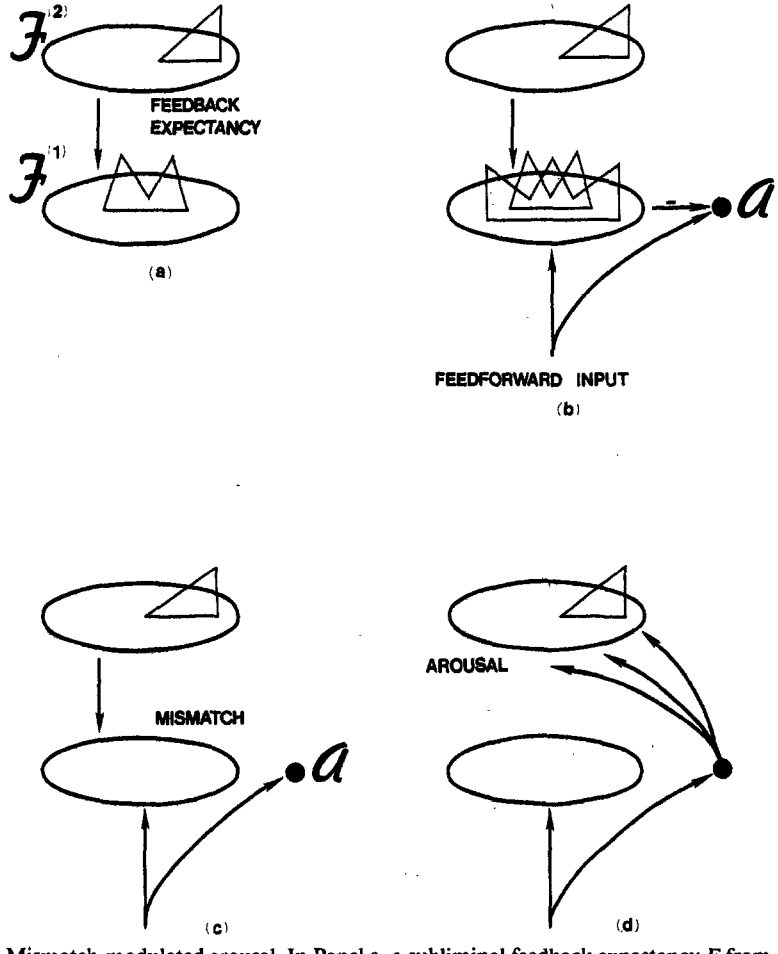

Figure 4. Mismatch-modulated arousal. In Panel a, a subliminal feedback expectancy E from $F^{(2)}$ to $F^{(1)}$ is maintained by short-term memory signaling from $F^{(2)}$. In Panel b, a feedforward input pattern U is also registered at $F^{(1)}$ as it simultaneously activates the arousal branch $\mathcal{A}$. In Panel c the mismatch between the two patterns across $F^{(1)}$ attenuates activity at this level. In Panel d inhibition at $F^{(1)}$ removes inhibition from $F^{(1)}$ to $\mathcal{A}$, thereby permitting $\mathcal{A}$ to unleash an arousal burst across $F^{(2)}$. Had the feedforward input matched the feedback expectancy, the amplified activity at $F^{(1)}$ would have inhibited the arousal source $\mathcal{A}$.

system $\mathscr{A}$ when $F^{(1)}$ is active. Because $\mathscr{A}$ also receives excitatory input due to the external event, the arousal subsystem $\mathscr{A}$ is disinhibited and releases a burst of nonspecific arousal across $F^{(2)}$ (Figure 4 [d]).

This description of the disinhibition of $\mathscr{A}$ assumes that an external event excites both the attentional and the orienting subsystems but that the orienting subsystem can only release signals when a mismatch within the attentional subsystem occurs. An early precursor of this idea is the Hebb (1955) hypothesis that every event has a cue and a vigilance, or arousal, function, A more recent correlate is my interpretation of $\mathscr{A}$ as the generator of the mismatch-negativity evoked potential (Näätänen, Hukkanen, & Järvilechto, 1980). A subtle aspect of this interpretation is that the mismatch occurs in the attentional subsystem, whereas the disinhibited mismatch negativity is elicited in the orienting subsystem.

17. STM Reset Versus STM Resonance

Before explaining how a pattern mismatch across $F^{(1)}$ can attenuate STM activity there, I need to discuss what happens if a pattern match occurs across $F^{(1)}$. By definition, such a match means that the external input that causes the feedforward pattern U is expected with respect to the feedback pattern E. The effect of this match is to amplify the STM activities of the matched pattern across $F^{(1)}$. Thus, a mismatch attenuates all activity across $F^{(1)}$, whereas a match amplifies patterned STM activity across $F^{(1)}$. I call the amplification of STM activity in a feedback module such as $F^{(1)}$ and $F^{(2)}$ an STM *resonance* (Grossberg, 1976b, 1978b, 1980).

Neurophysiological evidence for the existence of STM resonance has been described in the olfactory system by Freeman (1975, 1980, 1981) and in the visual system by Singer (1977, 1979). On a psychophysiological level, I identify this amplification of STM activity due to a pattern match with the *processing-negativity* or *match-detector* component of the N200 evoked potential complex (Grossberg, 1982a).

We have hereby been led to distinguish two functionally distinct actions of expected and unexpected events. Expected events can generate an STM resonance, whereas unexpected events can trigger selective STM reset. A subtle aspect of these complementary STM transactions is that both of them occur within the attentional subsystem, although only one of them is mediated by the orienting subsystem. In other words, the organization of the brain into *structurally* complementary subsystems does not correspond in an obvious way to the *functional* complementarity in the processing of expected and unexpected events.

A deeper subtlety of this functional interaction is implicit in the previous discussion and will be rendered explicit in Section 26. There I will conclude that the very STM resonance that represents paying attention to an event actively prepares the attentional subsystem to be reset by the orienting subsystem. An STM resonance does this by selectively depleting, or habituating, the transmitter gates in the active channels of the gated dipoles from which $F^{(2)}$ is constructed. If an arousal burst perturbs these dipoles, then STM will be rapidly reset.

18. Chunking and Expectancy Learning

The concepts of STM reset and STM resonance become fully meaningful when they are used to describe how the stability–plasticity dilemma is solved, notably how new internal representations (chunks) and expectancies are learned and remembered.

To complete this description, I assume that $F^{(1)}$ can send feedforward signals to $F^{(2)}$. The feedback expectancy signals from $F^{(2)}$ to $F^{(1)}$ are thus part of a reciprocal exchange of signals between successive levels in the network hierarchy. Anatomical correlates of this reciprocal exchange are the reciprocal thalamocortical connections that seem to occur in all thalamoneocortical systems (Macchi & Rinvik, 1976; Tsumoto, Creutzfeldt, & Legendy, 1976). Psychological correlates of this reciprocal exchange are the processes of recognition and recall that, when regulated by STM reset operations, lead to rapid hypothesis testing, or search, through associative memory (Grossberg, 1978a, 1978b).

The signals from $F^{(1)}$ to $F^{(2)}$ are gated by LTM traces before the signals reach their tar-

get cells. Then the gated signals are summed at each target cell, just as in Equation 10. The reader with some engineering background will recognize that the transformation from an output signal pattern $S = (S_1, S_2, \cdots S_m)$ to an input signal pattern $T = (T_1, T_2, \ldots, T_n)$, where $T_j = \sum_i S_i Z_{ij}$, defines a linear filter. Because the LTM traces Z_{ij} can change as a function of experience, this filter is called an *adaptive filter*. Both the adaptive filter due to $F^{(1)} \rightarrow F^{(2)}$ feedforward signaling and the learned expectation due to $F^{(2)} \rightarrow F^{(1)}$ feedback signaling obey the same gating laws and laws of associative learning. Their different intuitive interpretations are due to their different locations within the network as a whole.

When learning occurs in the LTM traces of the $F^{(1)} \rightarrow F^{(2)}$ pathways, the same input pattern U to $F^{(1)}$ will elicit a different STM pattern across $F^{(2)}$. Speaking intuitively, the internal representation of U across $F^{(2)}$ has changed. Just as an expectation is influenced by the $F^{(2)} \rightarrow F^{(1)}$ LTM traces, but is not identical with these traces, an internal representation is influenced by the $F^{(1)} \rightarrow F^{(2)}$ LTM traces, but is not identical with these traces.

19. *The Code of Adaptive Resonances*

With these comments in hand, I can now expand the notions of STM reset and STM resonance to include the self-organization process. If a feedback expectation E mismatches a feedforward input pattern U, then STM reset occurs before the LTM traces in the $F^{(1)} \rightarrow F^{(2)}$ pathways and the $F^{(2)} \rightarrow F^{(1)}$ pathways can change. Consequently, mismatched or erroneous interpretations of the environment cannot cause adaptive changes in LTM. The LTM traces can gate presently active signals prior to the reset event and can thereby alter network activity based on their past learning. However, the LTM traces are adaptively blind to present network activity, including the signals that they gate, because the LTM traces are slowly varying relative to the rapid time scale of filtering, mismatch, and reset.

By contrast, if a feedback expectation E matches a feedforward input pattern U, then STM is amplified across $F^{(1)}$. The signals from $F^{(1)}$ to $F^{(2)}$ are thereby amplified so that STM activity across $F^{(2)}$ is also amplified. Then the signals from $F^{(2)}$ to $F^{(1)}$ are amplified, and the entire network locks into an STM resonance. This global STM event defines the perceptual or attentive moment.

Resonant STM activities can be encoded in LTM because they are large enough and last long enough for the LTM traces to respond to them. An STM resonance is thus a context-sensitive global interpretation of the input data that signifies that the network as a whole considers this interpretation worthy of being adaptively incorporated into the network's memory structure. I call the dynamic process whereby LTM adapts to resonant STM patterns an *adaptive resonance.*

Using the notion of adaptive resonance, it is now easy to state what the perceptual or cognitive code of a network is, although the simplicity of this statement hides numerous subtleties and the need for much future scientific work. The code of a network is the set of stable adaptive resonances that it can support in response to a prescribed input environment.

James Gibson's lifelong ingenuity as a student of perceptual phenomena led him to conclude that the perceptual system "resonates to the invariant structure or is attuned to it" (Gibson, 1979, p. 249). Gibson has been criticized for emphasizing the phenomenal immediacy of perception at the cost of underemphasizing its processing substrates (Ullman, 1980). We can now understand that Gibson's emphasis was based on a correct intuition. The many processing steps, such as adaptive filtering, STM activation, readout of feedback expectancies, STM mismatch, disinhibition of orienting arousal, STM reset, and so forth, all involve perfectly good neural potentials, signals, and transmitters. However they are not accessible to consciousness. The conscious experience is, I suggest, the resonant or attentive moment, which seems to be immediate because it is a global event that energizes the system as a whole.

20. *The Noise-Saturation Dilemma*

Before turning to a consideration of conditioning data, we still need to understand

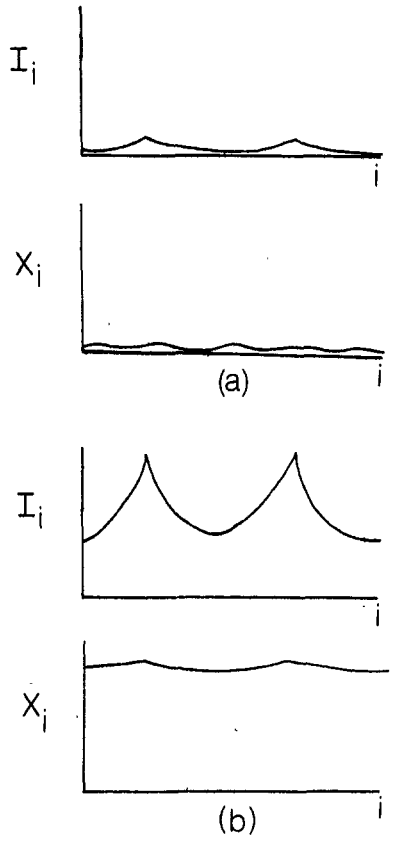

Figure 5. The noise-saturation dilemma. Panel a: At low input intensities, the input pattern $(I_1, I_2, \ldots, I_n)$ is poorly registered in the short-term memory activity pattern $(X_1, X_2, \ldots, X_n)$ because of internal noise in the cells. Panel b: At high input intensities, the input pattern is poorly registered in the short-term memory activity pattern because all of the cells' finitely many excitable sites get saturated.

how a pattern mismatch attenuates STM activity, how a pattern match amplifies STM activity, and how an STM resonance depletes the on-cell transmitter gate in a gated dipole. All of these properties follow from a study of competitive interactions between the cells at each of the several levels $F^{(1)}, F^{(2)}, \ldots, F^{(n)}$ of the network hierarchy.

The need for competitive interactions follows from a basic processing dilemma—*the noise-saturation dilemma*—that is faced by all cellular tissues, not only nerve cells, and must be solved before continuously fluctuating input patterns can be registered at all.

The noise-saturation dilemma is easy to state because it is so basic. Consider Figure 5. In Figure 5 (a), a differentiated pattern of inputs I_i is poorly registered in the activities, or STM traces, X_i because the overall input intensity is too small to override internal cellular noise. In Figure 5 (b), all the inputs are proportionally amplified to escape the cells' noisy range without destroying relative input importance (as when the reflectances of a picture remain constant despite its illumination at successively higher light intensities). Because all cells have only finitely many excitable sites, the smallest inputs are now large enough to turn on all the sites of their receptive cells; hence, the larger inputs might not be able to differentially excite their receptive cells, because there might be no more sites to turn on in these cells. The input differences are thereby lost because the activities saturate as all the cell sites are activated. Thus, in a cellular tissue, sensitivity loss can occur at both low and high input intensities. As the inputs fluctuate between these extremes, the possibility of accurately registering input patterns is imperiled.

I proved (Grossberg, 1973) that mass action competitive networks can automatically retune their sensitivity as inputs fluctuate to register input differences without noise or saturation contaminants. See Grossberg (1980, Appendices C and D) for a review. In a neural context these systems are called shunting on-center off-surround networks. Otherwise expressed, a network whose cells obey membrane equations (not additive equations) and that interact via an anatomy with a narrow excitatory focus and broadly distributed inhibition can automatically retune its sensitivity due to automatic gain control by the inhibitory signals.

21. STM Contrast Enhancement and Normalization: Hypothesis Testing and Overshadowing in a Limited-Capacity System

Because of the noise-saturation dilemma, competitive networks are ubiquitous wherever input patterns are accurately registered and transformed by cells. If these input patterns also need to be stored beyond the offset times of the inputs, as in STM representations that can remain active until later reinforcements can act upon them to influence

LTM encoding, then the competitive networks must also be feedback networks (Figure 6) whose positive feedback loops can reverberate the STM activities after the inputs cease.

Competitive feedback networks also exist for a deeper processing reason that is related to the stability–plasticity dilemma. These networks need to possess the properties of *contrast enhancement* and *normalization* to carry out hypothesis testing operations leading to the self-organization of new internal representations after an unexpected event occurs. The property of contrast enhancement is the capability of the network to attenuate small inputs and amplify large inputs before storing the contrast enhanced input pattern in STM. The property of normalization is the tendency of the total suprathreshold STM activity across the network to be conserved through time.

These properties enable the network to solve a problem that could prevent it from adaptively reacting to an unexpected event. Mismatch across $F^{(1)}$ causes an STM reset across $F^{(2)}$. If the nodes that are activated by $F^{(1)} \rightarrow F^{(2)}$ signals are hereby inhibited, then how does $F^{(2)}$ get activated at all by these signals during the next time interval? Why does the entire network not shut down? How does an unexpected input to $F^{(1)}$ ever get encoded if a mismatch shuts down the whole system?

The contrast-enhancement property supplies part of the answer. Not all of the nodes in $F^{(2)}$ that receive inputs from $F^{(1)}$ have their activities stored in STM. Only the nodes that receive relatively large inputs have their activities stored in STM. Only those nodes whose activities get stored in STM are reset by the mismatch-modulated arousal burst. The nodes that receive smaller inputs are not reset because they did not reverberate in STM. These latter nodes can still respond to the signals from $F^{(1)}$ to $F^{(2)}$ in the next time interval.

The contrast-enhancement property thus shows how some input-activated nodes can be spared by the STM reset process. However, this property is not sufficient because these nodes, after all, receive such small inputs that they could not previously reverberate in STM. Why can they reverberate in STM after the nodes that received large inputs are inhibited by dipole rebounds?

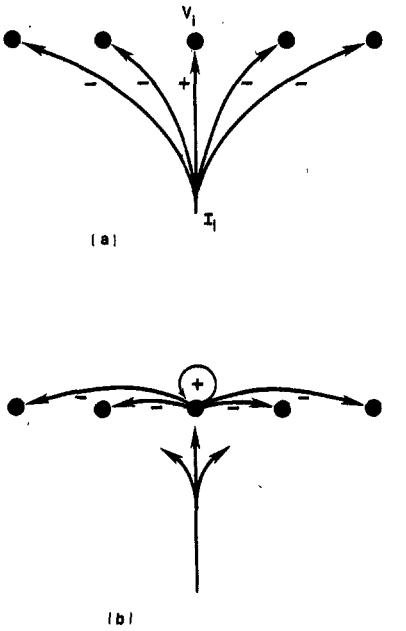

Figure 6. Two types of competitive networks. Panel a: A feedforward competitive network delivers both excitatory and inhibitory inputs to its receptive cells. This competitive input distribution allows the cells to respond to input patterns of widely varying background intensity without saturation. Panel b: A feedback competitive network generates excitatory and inhibitory feedback signals among its own cells. When these feedback signals are sigmoid, or *S*-shaped, functions of cell activity, the network can contrast-enhance the input pattern before storing it in short-term memory. This contrast-enhancement property follows from the sigmoid signal's ability to suppress, rather than amplify, noise through the network's excitatory feedback loops. The network also tends to conserve its total suprathreshold activity through time. This normalization property dynamically explains the limited capacity of short-term memory as a consequence of solving the noise-saturation dilemma.

The normalization property now comes to the rescue. The total suprathreshold STM activity tends to be conserved. Because the dipole-inhibited nodes can no longer compete for this conserved activity, the remaining input-excited nodes inherit it (Figure 7). Thus, the nodes that fared poorly in the competition for STM activity in the original field $F^{(2)}$ fare much better in the "renormalized" field wherein their strongest competitors have been inhibited by dipole rebounds. The successive renormalization of the field $F^{(2)}$ by rapid reset events can be viewed as a type

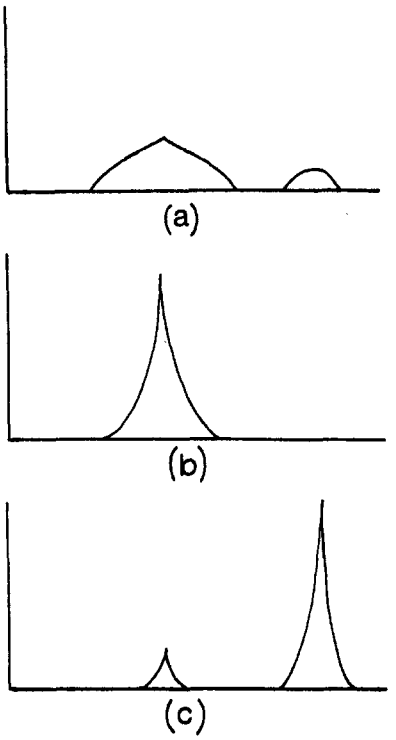

Figure 7. Renormalization. The input pattern in Panel a elicits short-term memory activity pattern $X_i^{(2)}$ in Panel b by suppressing small inputs and contrast-enhancing large inputs. After $X_i^{(2)}$ is suppressed by dipole rebounds, the small input activities inherit normalized activity from the suppressed populations to elicit the distinct short-term memory activity pattern in Panel c.

of parallel hypothesis testing, or principal-component analysis, whereby the field zooms in on nodes capable of generating an STM resonance.

The way in which STM normalization influences the processing of unexpected events is relevant to my analysis of overshadowing. Because of STM normalization, increasing the STM activity of one representation forces a decrease in the STM activities of the representations with which the enhanced representation competes. Otherwise expressed, STM normalization provides a dynamical explanation of why STM is a limited capacity system.

It is important to realize that although competitive feedback networks possess the normalization property, competitive feedforward networks do not. An input is small in a feedforward network because it has already fared badly in the feedforward competitive interaction. The normalization property must occur after the stage of input delivery, not before or at this stage.

22. Overshadowing and Sigmoid Signal Functions

The discussion of contrast enhancement and normalization indicates some properties that are needed to accomplish hypothesis testing in response to an unexpected event. This discussion does not, however, show how these properties are obtained or whether other important properties coexist with them. We now need to review this issue in a more detailed fashion, because an important property that we need is still not available.

The need is clarified by considering a standard overshadowing experiment (Kamin, 1968, 1969). Suppose that an animal receives a series of CER conditioning trials wherein a Cue A (such as a light) is followed by a standard shock US. Let these trials be followed by a series of CER conditioning trials to the compound cue AB (such as a light–tone combination) followed by the same shock US. A later test of Cue B's ability to suppress bar pressing for food shows that its fearfulness has been blocked by the prior conditioning of Cue A to the shock. By contrast, if the compound conditioning trials use a different level of shock as a US than was used during the conditioning of Cue A, then the conditionability of Cue B is restored.

In the light of the previous theoretical discussion, one might wish to say that the unexpected change of US at the onset of compound trials causes an STM reset that somehow amplifies Cue B's STM representation and thereby frees it from blocking. This wish is, however, incompatible with Equation 9, which shows that all active representations are reset if any representation is reset. Is the present theory incompatible with basic facts of overshadowing?

At this point one must soberly face the fact that the theory would collapse without the benefit of mathematics. Indeed, it is quite impossible to understand the STM transformations during conditioning experiments without a suitable mathematical tool.

Equation 9 was derived under the hypothesis that a linear signal function transmutes

the arousal signal I and the test signal J into signals S_1 and S_2 to the dipole gates. I will now say why a linear signal function can never be used in a competitive feedback network within a perceptual or cognitive processor. In fact, a linear signal function does not imply the contrast-enhancement property that is needed to reset STM in response to an unexpected event. The right kind of signal function gives all the properties that are needed, including the overshadowing property. This signal function is a sigmoid, or S-shaped, signal function. Such a signal function acts like a threshold at low cell activities and levels off at high cell activities.

23. Noise Suppression and Selective Enhancement, or Dishabituation, of Overshadowed Representations by Unexpected Events

A sigmoid signal function is needed for a basic processing reason. The positive feedback signaling in competitive feedback networks can be a mixed blessing if it is not properly designed. Positive feedback can subserve such desirable properties as STM storage and normalization. It can also, if improperly designed, flood the network with noise generated by its own activity. This noise amplification property will have disastrous effects on network processing whenever it occurs, and it can occur in formal network "seizures" and "hallucinations" (Ellias & Grossberg, 1975; Grossberg, 1973; Kaczmarek & Babloyantz, 1977; Schwartz, 1980).

A proper choice of signal function prevents noise amplification (Grossberg, 1973). The simplest physically plausible signal function that prevents noise amplification is the sigmoid signal function. The opposite of noise amplification is noise suppression. This noise suppression property attenuates small inputs to the network. By normalization, it thereby amplifies large inputs to the network. A competitive feedback network's contrast enhancement property is thus a variant of its noise suppression capability.

We can now reanalyze the response of a gated dipole to an arousal burst when a sigmoid signal, as opposed to a linear signal, is used. When this is done, one finds that the same arousal burst that rebounds the on-responses of very active dipoles will *enhance* the on-responses of weakly active dipoles. In other words, overshadowed representations in a dipole field can actually be enhanced by the same surprising event that inhibits more salient representations via antagonistic rebound. That all the properties that are needed occur automatically as a result of basic processing constraints like noise suppression is what I call a minor mathematical miracle. Minor mathematical miracles should not be taken lightly. They usually mean that the intuitive ideas that they reify contain a lot of truth.

To illustrate this new rebound property, let the sigmoid signal function be $f(w) = w^2(1 + w^2)^{-1}$. Because we are interested in the smallest, or threshold, increment ΔI that can cause a rebound, we can approximate $f(w)$ by w^2. Then, in the gated dipole of Figure 1, a rebound occurs to an arousal increment ΔI given a previous arousal level I and on-input J only if

$$\Delta I > g(I, J), \tag{12}$$

where the function $g(I, J)$ is defined by

$$g(I, J) = \frac{A - I(I + J) + (A + I^2)^{1/2}[A + (I + J)^2]^{1/2}}{2I + J}. \tag{13}$$

Because $g(I, J)$ is a decreasing function of J, it is easier to rebound an intensely activated dipole ($J \gg 0$) than a weakly activated dipole ($J \cong 0$).

If $\Delta I < g(I, J)$, then the arousal increment ΔI can cause an enhanced on-reaction rather than an off-rebound. These properties illustrate how the habituation of cell responses due to transmitter depletion can be dishabituated by an unexpected event. An electrode that is too coarse to distinguish between on-enhancements and off-rebounds near its recording site could record dishabituation of this type in response to an arousal increment at all electrode placements near previously active cells, whereas a more sensitive elec-

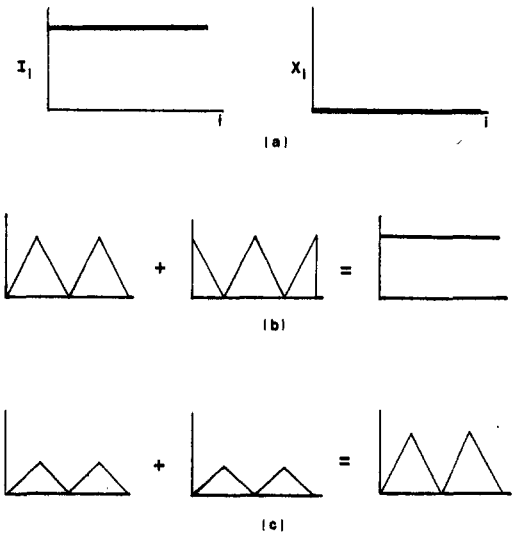

Figure 8. Pattern matching. In Panel a the noise suppression property converts a uniform (or zero spatial frequency) input pattern (I_i) into a zero activity pattern (X_i). In Panel b two mismatched input patterns add to generate an approximately uniform input pattern, which is suppressed by the mechanism of Panel a. In Panel c two matched patterns add to yield a total input pattern that is more active than either pattern taken separately.

trode could record both on-enhancements and off-rebounds as dishabituation reactions at some cell locations, as well as off-reactions at other cell locations.

24. Noise Suppression and Pattern Matching

A variant on the noise suppression theme implies a mechanism of pattern matching at $F^{(1)}$. A shunting competitive network can easily be designed to suppress uniform input patterns (Figure 8 [a]). Such patterns represent noise in the sense that no node is differentiated from any other by the input. If a network can suppress a uniform input pattern, then it can suppress a sum of two mismatched input patterns (Figure 8 [b]) because the mismatched peaks and troughs of the input patterns add to produce an almost uniform total pattern. By contrast, a sum of two matched input patterns yields an amplified network reaction (Figure 8 [c]) because the network's shunting mechanism reacts to the sum of two patterns with a larger gain than to a single input pattern. See Grossberg (1980 or 1981a) for mathematical details.

25. Sigmoid Signals and Tuning: Multiple Effects of a US

One final consequence of sigmoidal signaling in a competitive feedback network can help us to understand conditioning experiments. Noise suppression due to sigmoid signaling also implies the following interesting property. A competitive feedback network undergoing sigmoidal feedback signaling possesses a parameter called a *quenching threshold* (QT). The QT defines the noise level of such a network, because the network inhibits activities that start out less than the QT and contrast-enhances activities that exceed the QT before storing them in STM. Any network with a QT can be tuned; by varying the QT the network's ability to store or suppress inputs can be altered through time.

For example, if an arousal source nonspecifically increases the level of shunting inhibition acting on a competitive network's feedback inhibitory interneurons, then the net disinhibitory action will cause the network's QT to decrease. The network's STM storage of input patterns will thereby be facilitated. This type of arousal event should not be confused with orienting arousal. It is the type of arousal that lowers and heightens the sensitivity of the attentional subsystem, as during sleep and wakefulness.

I will now briefly indicate how the concept of an attentional QT can be implicated during a conditioning experiment. To do this I will use some concepts intuitively that will be precisely defined in the next sections. Suppose that any cue that can activate a drive representation can also reduce the attentional QT. Then an unexpected US can have three distinct effects on attentional processing. As an unexpected event that mismatches active feedback expectancies, the US can remove some STM representations from overshadowing by differentially amplifying them. As a US per se, it can activate a drive representation and thereby further abet the STM storage of overshadowed cues by lowering the QT. This effect sensitizes the processing of all cues that survive STM reset and occurs

later in time than STM reset. Finally, as a generator of conditioned incentive motivational feedback, the US can differentially strengthen STM activities of motivationally compatible cues. By contrast, an unexpected CS is only capable of eliciting differential STM reset because of mismatch with feedback expectancies.

All of the three effects itemized above depend on learning, but in different neural pathways and on different time scales. The first effect involves expectancy learning, the second effect involves conditioned reinforcer learning (which enables a cue to turn on a drive representation), and the third effect involves incentive motivational learning. If such a QT effect of a US exists, its generator will probably be turned on by signals that are elaborated within the hippocampus, which I identify as the final common pathway of the drive representations (Grossberg, 1971, 1975, 1980).

Conditioning and Attention

26. *Gated Feedback Dipoles*

We now have all the conceptual threads that we need to discuss conditioning and attention. Some loose threads still need to be tied, however. Then we will find that a large body of data falls into place with little difficulty. Because a deep theoretical understanding comes from knowing how particular mechanisms generate particular constellations of data properties, I will build up the mechanisms in stages and will supply data markers at each stage.

Two distinct design principles coexist at the same cells: gated dipoles and shunting competitive feedback networks. I will now show how to join these two principles into a single network. This can be done by making explicit a property that was mentioned in Section 22 without pursuing its implications. There I suggested that a nonspecific arousal burst could reset STM even after the cues that initiated STM storage had terminated. I used this property to begin an explanation of overshadowing. In order for this to happen, the STM feedback loops must contain the transmitter gates so that STM activity can differentially deplete the active STM loops and thereby prepare them for rebound. Figure 9 summarizes this structural arrangement by depicting a conditionable input pathway abutting a gated feedback pathway. I will henceforth assume that the transmitter in a conditionable input pathway is cholinergic and that the transmitter in a gated STM feedback loop is catecholaminergic (Grossberg, 1972b), because a large body of data is compatible with this suggestion (Butcher, 1978; Epstein, Kissileff, & Stellar, 1973; Friedhoff, 1975a, 1975b).

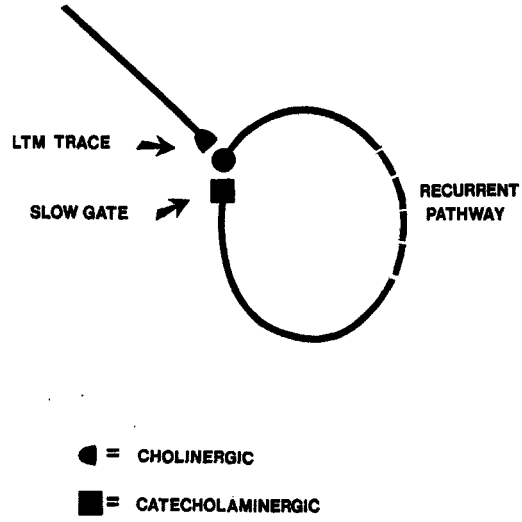

Figure 9. Slow gated feedback and long-term memory interaction. A conditionable pathway that is excited by a cue's internal representation feeds into a gated feedback loop. Because the gate occurs within a feedback loop, the conditionable pathway can achieve two distinct processing objectives. It can sample and store in its long-term memory (LTM) trace the gated output of the loop. It can also alter the short-term memory activity in the loop by changing its own signals through time. The LTM trace in the conditionable pathway is assumed to be part of a cholinergic interaction. The gate in the feedback loop is assumed to be part of a catecholaminergic interaction. The text and Figure 10 show how these components can be embedded into a gated feedback dipole.

There is another way to derive Figure 9 even when no arousal-initiated rebound exists. This alternative derivation holds when the off-cells represent features or behavioral categories that are complementary to those represented by the on-cells: for example, fear versus relief, hunger versus satiety, vertical red bar on green field versus vertical green bar on red field, and so forth. The derivation proceeds in three steps (Grossberg, 1972b).

1. Sudden offset of a conditioned cue input can cause a rebound, much as offset of a conditioned source of fear can elicit relief (Denny, 1971; Masterson, 1970; McAllister & McAllister, 1971). To accomplish this rebound the cue input is delivered to the network at a stage *before* the transmitter gate. Only in this way can the cue deplete the gate so that its offset can drive a rebound.

2. Onset of a cue input can elicit sampling signals capable of encoding a rebound in LTM, much as a tone that is turned on contingent on shock offset can become a source of conditioned relief (Dunham, 1971; Dunham, Mariner, & Adams, 1969; Hammond, 1968; Rescorla, 1969; Rescorla & LoLordo, 1965; Weisman & Litner, 1969). Thus, the cue input is delivered to the network at a stage *after* the transmitter gate, where the rebound can be sampled.

3. Properties 1 and 2 are true for *all* cues that can be conditioned to these categories, because whether a given cue will be conditioned to the onset or to the offset of any particular category is not known a priori. Every cue input is delivered both before and after the transmitter gating stage. The transmitter gate thus occurs in a feedback pathway, as in Figure 9.

The existence of two distinct derivations leading to a similar network design is important, because not every recurrent network that can be reset by offset of a cue need possess a mismatch-contingent arousal source, even though an arousal source per se is required. These derivations suggest that the anatomical design in Figure 9 is basic and that the input mechanisms that control rebound in this common design can be adapted to satisfy specialized processing constraints.

One further constraint can readily be satisfied by this design. The cue inputs arrive before the stage of dipole competition so that at most one of the complementary outputs (on-cell versus off-cell) can be positive at any time. The next section depicts the minimal anatomy that joins together gated dipole feedback pathways and conditionable cue input pathways that terminate before the dipole competition stage. This microcircuit is needed to build up both the sensory and the drive representation networks that are activated by conditioning and attentional manipulations.

27. Drives, Conditioned Reinforcers, Incentive Motivation, and CNV

Figure 10 depicts such a minimal anatomy and assigns to its pathways the motivational interpretation that turns the anatomy into a drive representation. In Figure 10 the specific inputs to the gated dipole are of two kinds: internal drive inputs and external cue inputs. In the case of eating, the positive drive input increases with hunger, and the negative drive input increases with satiety, owing either to gastric distension or to slower metabolic factors (Anand & Pillai, 1967; Janowitz, Hanson, & Grossman, 1949; Le Magnen, 1972; Sharma, Anand, Dua, & Singh, 1961). Let the drive inputs and the nonspecific arousal input be gated by a catecholaminergic transmitter in both the on-channel and the off-channel to calibrate the proper relative sizes of dipole on-responses and off-responses.

Let each external cue input send a conditionable pathway to both the on-channel and the off-channel of the dipole. Each cue can become a conditioned reinforcer of either positive or negative sign, depending on which of its LTM traces in the on-channel or the off-channel is larger. To calibrate the relative sizes of these LTM traces in an unbiased fashion, I assume that the transmitter system that subserves LTM encoding in both branches is the same and is cholinergic. These chemical interpretations may eventually have to be changed, but the processing requirements of accurately calibrating relative rebound or conditioned reinforcer strength across competing channels are robust.

The cells at which external cue, drive, and arousal inputs converge are assumed to be *polyvalent*: These cells only fire vigorously when both their external cue and their internal drive inputs are sufficiently large. The outputs from these polyvalent cells compete before generating a net dipole output in either the on-cell or the off-cell channel, but not both. These dipole outputs play the role of *incentive motivation* in the theory. The existence of polyvalent cells is compatible with the intuition that incentive motivation need not be released even when drive is high if compatible cues are unavailable. For example, a male animal left alone will busily do the many things characteristic of his spe-

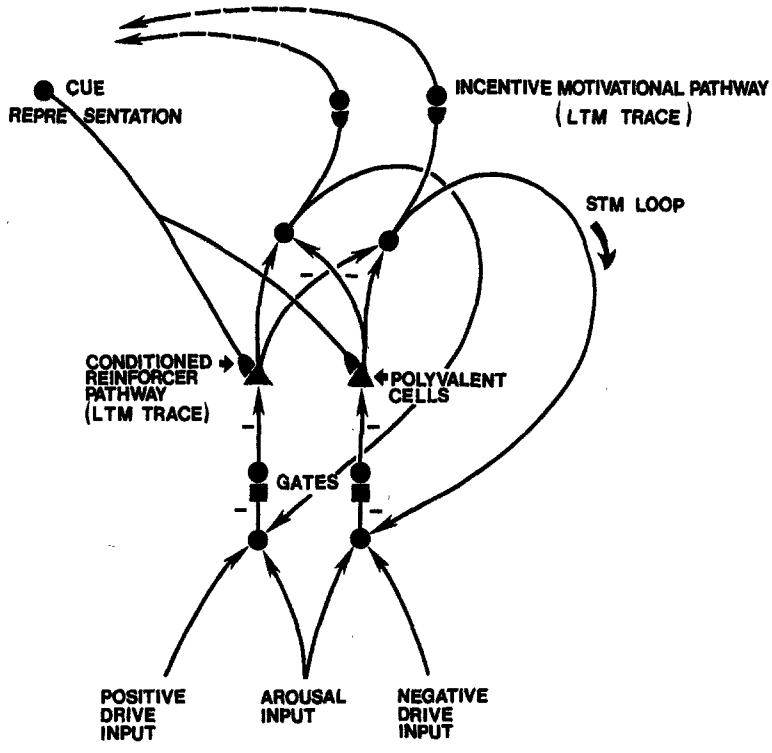

Figure 10. A motivational network: network with a gated feedback dipole hooked up to conditionable pathways from the internal representations of cues. The text describes how the feedback loops between external cue representations and internal drive representations, and between internal drive representations and themselves, join together mechanisms of reinforcement, drive, incentive motivation, competition, arousal, and short-term memory (STM). LTM = long term memory.

cies, 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). Incentive motivation need not be released even when compatible cues are available if drive is low. Seward and his colleagues (Seward & Proctor, 1960; Seward, Shea, & Elkind, 1958; Seward, Shea, & Davenport, 1960) found that if an animal is not hungry, then no amount of food will be adequate to reinforce its behavior.

Figure 10 suggests a different relationship between drive and incentive than is found in Routtenberg (1968). Routtenberg makes incentive and drive the complementary concepts of his two-arousal hypothesis. In Figure 10 drive and incentive are both part of the attentional subsystem, which is complementary to the orienting subsystem. In my theory drive and arousal are not the same concept.

28. Extinction, Conditioned Emotional Responses, Conditioned Avoidance Responses, and Secondary Conditioning

To indicate how a gated feedback dipole works, the next two sections summarize some of its formal properties using reinforcement and motivation terminology. This section reviews how each cue's LTM sampling of both the on-channel and the off-channel contributes to the explanation of some basic conditioning processes (Grossberg, 1972a, 1972b) and thereby indirectly supports the claim that both LTM pathways need to be built up from similarly calibrated transmitter mechanisms. The simplest LTM law says that LTM encoding occurs only when a cue pathway and a contiguous polyvalent cell are simultaneously active (Grossberg, 1964, 1968; Hebb, 1949).

Suppose that during early learning trials, a cue occurs just before a large unconditional

signal, such as a shock, turns on the on-channel. This unconditional signal elicits the incentive output of the on-channel, which triggers a fear reaction, even before learning occurs. By associating the cue with shock on learning trials, its LTM trace abutting the on-channel grows much larger than its LTM trace abutting the off-channel. If the cue is then presented by itself, the LTM-gated signal to the on-channel is much larger than the LTM-gated signal to the off-channel. The on-channel wins the dipole competition, so the cue elicits fear. The cue has to thereby become a conditioned reinforcer that can elicit a conditioned emotional response, or CER (Estes, 1969; Estes & Skinner, 1941).

Now suppose that after the cue has become an elicitor of a CER, environmental contingencies change. Suppose that the cue no longer reliably predicts future events and that an unexpected event occurs while the cue is on. Suppose that the unexpected event triggers an antagonistic rebound in the off-channel. Because the cue is on, its LTM trace abutting the off-channel will grow. If this occurs sufficiently often, the off-LTM trace will grow as large as the on-LTM trace. After this happens, presenting the cue will generate comparable LTM-gated signals to both the on-channel and the off-channel. After these signals compete, the net incentive motivation output will be very small. The cue no longer elicits a CER. It has been rapidly extinguished by unexpected events.

This cue is extinguished because it remains on both before and after the unexpected event. It is an irrelevant cue with respect to the contingency that triggered the unexpected event. By contrast, a cue that turns on right after the unexpected event occurs will sample only the off-reaction of the dipole. Only its LTM trace abutting the off-channel will grow. Later presentation of the cue will thereby elicit a large off-reaction. If the off-reaction corresponds to a relief reaction, then the cue has become a source of conditioned relief by being paired with offset of a source of fear. Although the cue has never been paired with a positive reward, it thereafter can be used as a positive reinforcer or source of consummatory motivation. This mechanism helps us to understand how avoidance behavior can be persistently maintained long after an animal no longer experiences the fear that originally motivated the avoidance learning (Grossberg, 1972a, 1972b, 1975; Maier, Seligman, & Solomon, 1969; Seligman & Johnston, 1973; Solomon, Kamin, & Wynne, 1953).

A similar argument shows how secondary conditioning can occur. For example, offset of a positive (or negative) conditioned reinforcer S_1 can drive an antagonistic rebound that conditions a cue S_2 whose onset is contingent upon the offset event to be a negative (or positive) conditioned reinforcer. This mechanism uses the feedback in the gated dipole in a major way. Offset of the reinforcer S_1 can elicit a rebound because it occurs at a stage prior to the gate, whereas sampling of the rebound can occur because Cue S_2 delivers it signals at a stage subsequent to the gate.

29. Motivational Baseline, Switching, and Hysteresis

The positive feedback loops in the gated dipole of Figure 10 turn this network into a feedback competitive network. The slowly varying transmitter gates do not alter the fact that a feedback gated dipole shares many properties with other feedback competitive networks. For example, the dipole now has an STM storage capability, which means that it can defend a motivational decision against sufficiently small momentary fluctuations in cue or arousal inputs. This hysteresis property helps to provide the inertia needed to carry out a sustained, motivated act during irrelevant environment perturbations. The STM normalization property refines this capability by maintaining a temporally stable baseline of incentive motivation. The contrast-enhancement property due to sigmoidal feedback signaling helps to control sharp motivational switching between behavioral alternatives when the net balance of inputs succeeds in overcoming the hysteretic inertia of the previous motivational choice.

Frey and Sears (1978) have suggested a formal model to explain some of these properties. Their model builds sigmoid and hysteresis properties into a cusp catastrophe model of conditioning and attention. Although their model provides one way to vi-

sualize sudden switches, the catastrophe variables do not correspond to physical variables, and the model provides no physical explanation of why sigmoid and hysteresis properties appear in the data. The gated dipole theory provides a physical explanation that does not correspond to a cusp catastrophe, and it implies a large body of data and predictions that are invisible to the cusp picture. For example, when sigmoid feedback signals are used in a gated dipole, this network also possesses inverted U properties that are reflected in a large body of data about normal and abnormal behavior (Grossberg, 1972b, 1982a). These inverted U properties are part of the minor mathematical miracle.

Because of the importance of the gated dipole concept to my motivational theory, I will also discuss how the same mechanisms work in the case of hunger and satiety, rather than shock. In the case of hunger, a positive drive input increases with hunger, whereas a negative drive input increases with satiety. An increase in the positive drive input disinhibits the polyvalent cell that is two inhibitory synapses away. Suppose that this polyvalent cell also receives a large conditioned signal from a cue representation. In other words, the cue is a conditioned reinforcer with respect to this drive representation, and the cue is active in STM. Then the polyvalent cell can vigorously fire. Suppose at this moment that the negative drive input is small (e.g., the hunger level is high) and/or only weak conditioned signals reach the polyvalent cell of the negative drive representation. Then this polyvalent cell does not fire vigorously. Thus, after competition takes place, the positive drive channel wins. It can therefore emit incentive motivation signals to the cue representations. These conditionable signals help to regulate attention by modifying the total excitatory input pattern to the cue representations.

The positive drive channel can also deliver excitatory feedback to the cells that receive positive drive input. This excitatory feedback can sustain the activity of the positive drive representation. It can thereby store a motivational decision in STM against small input perturbations (hysteresis), maintain a steady motivational baseline (normalization), and regulate sharp motivational switching (contrast enhancement). All of these properties are STM properties. The sustained STM reverberation also allows contiguous LTM traces of active cue representations to encode the large activity of the positive drive representation at a rate that increases with the STM activity of the cue representation. These active cues can thereby become positive conditioned reinforcers.

If the incentive motivation from the positive drive representation supports sustained, motivated behavior (e.g., eating), then the negative drive input slowly grows (e.g., satiety increases). The increase in the negative drive input shuts off STM at the positive drive representation via the competitive interaction. The motivated behavior thereby loses its incentive motivation support (e.g., eating stops).

If positive conditioned reinforcer input is rapidly withdrawn before the negative drive input increases, then an antagonistic rebound can be elicited in the negative drive channel. This rebound rapidly terminates the motivated behavior. An antagonistic rebound can occur because a sudden reduction of positive conditioned reinforcer input reduces the signal within the feedback loop of the positive drive representation. The total signal to the transmitter gate in the positive drive channel is thereby reduced, and a rebound is elicited just as in Figures 1 and 2. A cue whose LTM traces sample the antagonistic rebound can become a negative conditioned reinforcer. A cue whose LTM traces sample both the positive drive representation and the negative drive representation is extinguished (irrelevant) with respect to this drive, because its positive and negative gated signals to the gated dipole inhibit each other at the competitive stage before any net incentive motivation can be released.

An issue of some importance concerns how strict the polyvalent constraint is on the firing of cells where external cue inputs and drive inputs converge. To illustrate the issue, suppose that a satiety input grows because of sustained eating. If the polyvalent constraint is strict, then the polyvalent cell that receives the large satiety input cannot fire at all unless it also receives a large cue input. If not, the satiety input cannot inhibit the positive incentive motivation that was sup-

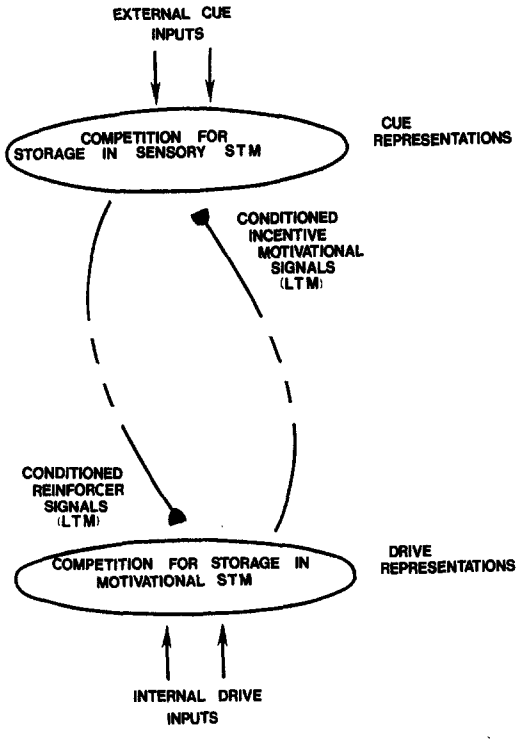

Figure 11. Adaptive resonance between dipole fields. When external cues excite the short-term memory (STM) traces of their internal representations, these internal representations elicit signals that are distributed nonspecifically across the various internal drive representations. During conditioning the pattern of reinforcing and drive inputs to the drive representations can alter the long-term memory (LTM) traces within certain of these signal pathways, as Figure 10 has illustrated. The corresponding external cues thereby acquire conditioned reinforcer properties. On recall trials the conditioned reinforcer signals from external cues combine with internal drive inputs at the drive representations to determine which drive representations will fire. Output from the drive representations plays the role of incentive motivation in the network. Incentive motivation is released from a given drive representation only if the momentary balance of conditioned reinforcer signals plus drive inputs competes favorably against these factors within the other drive representations. The incentive motivational pathways are also nonspecific and conditionable. External cue representations that receive large incentive motivational feedback signals are favored in the competition for storage in sensory short-term memory, as Figure 13 describes in greater detail.

porting eating. If the polyvalent constraint is not strict, then small background cue inputs will suffice to permit polyvalent cell firing in response to sufficiently large satiety inputs. This problem is overcome by the network of Figure 12, because the drive inputs compete in that network before they influence the polyvalent cells. Once the concept of polyvalence is before us, we can begin to classify which networks best marry polyvalence to other important processing constraints.

30. Adaptive Resonance Between Dipole Fields

Now that we have a clearer view of how to design the microcircuitry of a gated feedback dipole, we need to build these microcircuits into a global processing scheme. The gated feedback dipoles are part of dipole fields wherein on-cells are joined by shunting competitive feedback networks, off-cells are joined by shunting competitive feedback networks, and on-cells are joined to off-cells via gated dipoles. The dipole fields themselves interact via adaptive filters, much as the fields $F^{(1)}$ and $F^{(2)}$ interact in Section 18.

Figure 11 depicts the macrocircuit that will be most prominent in my discussion of conditioning data. It describes a feedback module wherein sensory and drive representations send signals to each other via nonspecific excitatory conditionable pathways (adaptive filters). These representations are organized into dipole fields. Each dipole field is capable of STM contrast enhancement, normalization, hysteresis, and rebound. The interfield conditionable pathways send branches to both the on-cells and the off-cells of the dipole fields, just as they do to explain extinction and secondary conditioning in Section 28.

The conditionable pathways from sensory representations to drive representations encode the conditioned reinforcer properties of external cues. The conditionable pathways from drive representations to sensory representations encode the incentive motivation properties of internal drives. An adaptive resonance occurs within this network when the reinforcing properties of active external cues sufficiently match the motivational properties of active internal drives to lock STM into a global interpretation of the data.

In the theory that I developed in Grossberg (1971, 1972a, 1972b, 1975), the final processing stage in the external cue representations is assumed to be cortical, and the final

processing stage in the drive representations is assumed to be hippocampal. Gabriel, Foster, Orona, Saltwick, and Stanton (1980) summarize recent data that support a qualitatively similar conclusion. They write, "the hippocampal formation is a region critical for encoding or 'modelling' of stimulus-reinforcement contingencies" (p. 189). They note that the hippocampus is reciprocally connected with cingulate cortex and with the anteroventral nucleus of the thalamus and summarize data suggesting that cortical "codes inappropriate to the stimulus item being presented would create mismatch with the hippocampal model, thereby eliciting code-suppression in cortex and thalamus" (p. 189).

31. A Motivational Dipole Field: Drive–Reinforcer Matching and Motivational Competition

Now we need to fill in the microcircuitry of the dipole fields using Section 27 as a guide. Figure 12 depicts an anatomy that possesses the minimal drive representation properties that I will need. In this anatomy each motivational channel possesses a posi-

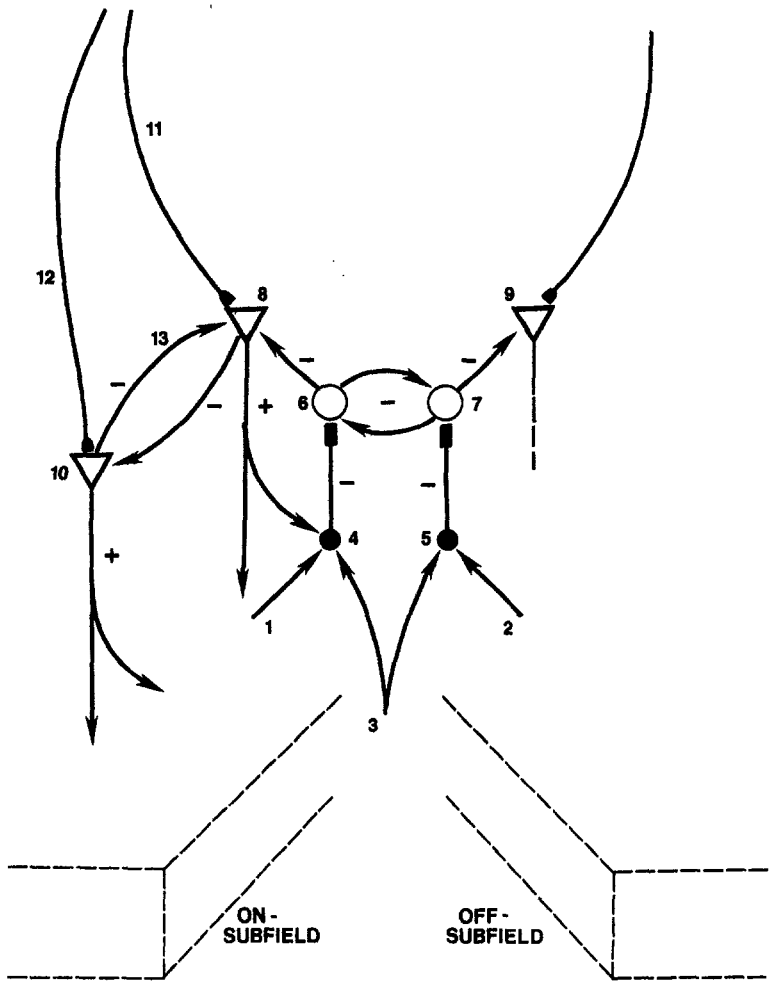

Figure 12. A motivational dipole field. The text describes how individual motivational dipoles are joined together by competitive feedback networks to regulate which dipole(s) will reverberate in STM and thereby release positive or negative incentive motivation in response to the changing balance of drive inputs, arousal inputs, and conditioned reinforcing signals.

tive feedback loop that runs through a gated dipole. These positive feedback loops are the on-centers of a competitive feedback network that joins together motivational channels. The competitive feedback network provides a matching interface that runs across the motivational channels. At this interface, spatial patterns of (conditioned) reinforcer signals are matched with spatial patterns of drive signals. Only a sufficiently good match can trigger sustained polyvalent cell firing (Sections 17 and 27). If this network is tuned (Section 25) so that only a single winning channel can reverberate in STM, then sharp motivational switching will occur. Such a setting of the network defines the channels as being motivationally incompatible. A lower setting of the QT can permit compatible combinations of drive representations to be synergistically activated. Possible settings depend on the choice of numerical network parameters and can vary across species and individuals without changing the underlying design principle.

To understand in greater detail how a motivational dipole field works, I will summarize the processing stages in Figure 12 step by step. Pathways 1 and 2 carry specific, but complementary, drive inputs (e.g., hunger vs. satiety) to a single dipole. Pathways labeled 3 carry nonspecific arousal to this dipole. Cells 4 and 5 add these inputs and thereupon inhibit the tonically active Cells 6 and 7. (Tonic cells have open symbols; phasic cells have closed symbols.) Pathways $4 \rightarrow 6$ and $5 \rightarrow 7$ contain slow transmitter gates (square synapses), assumed to be catecholaminergic. If Input 1 exceeds Input 2, then the transmitter in Pathway $4 \rightarrow 6$ is depleted more than the transmitter in Pathway $5 \rightarrow 7$, thereupon 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 imbalance disinhibits tonic Cell 8 and further inhibits tonic Cell 9. Both Cells 8 and 9 are polyvalent, meaning that all their excitatory inputs must be active for these cells to vigorously fire. (Triangles denote polyvalence.) The polyvalent cells are assumed to be pyramidal cells. Because Cells 8 and 9 are polyvalent, a larger input to Cell 1 than Cell 2 cannot fire these cells. However, such an imbalance can prevent Cell 9 from firing.

To see how Cell 8 can fire, we consider the polyvalent cells, 8 and 10, of two different motivational channels. Cells 8 and 10 compete via the inhibitory (interneuronal) Pathways 13. The polyvalent Cells 8 and 10 also receive inputs from external cue representations via the conditionable Pathways 11 and 12, respectively, whose LTM traces (within the filled hemicircles abutting Cells 8 and 10) encode conditioned reinforcer properties of their respective external cues. These LTM traces are assumed to be cholinergic.

The conditioned reinforcer inputs combine with drive and arousal inputs at their respective polyvalent cells, which begin to fire if their thresholds are exceeded. The polyvalent cells thereupon compete among themselves via the "intrinsic" feedback inhibitory Pathways 13, as they simultaneously try to excite themselves via positive feedback pathways such as $8 \rightarrow 4 \rightarrow 6 \rightarrow 8$.

If, for example, Cell 8 wins this competition, then the transmitter gate in Pathway $4 \rightarrow 6$ is depleted as the suprathreshold reverberation bursting through Cell 8 via Pathway $8 \rightarrow 4 \rightarrow 6 \rightarrow 8$ drives LTM changes in Pathway 11. The reverberation thus induces conditioned reinforcer changes even as it prepares the network for motivational reset by rapid offset of Pathway 11 or a rapid increment in Pathway 3.

32. A Sensory Dipole Field: The Synchronization Problem and DC Potential Shifts

Figure 13 depicts the minimal anatomy that I will need to join together external cue representations. This dipole field possesses additional structure compared to Figure 12 because it solves a specialized design problem, which I call the *synchronization problem* of classical conditioning. The synchronization problem recognizes that without specialized network buffers, Pavlovian associations could rapidly extinguish whenever a CS and US were presented with different interstimulus delays on successive learning trials. The synchronization problem was solved in Grossberg (1971) and provided the impetus

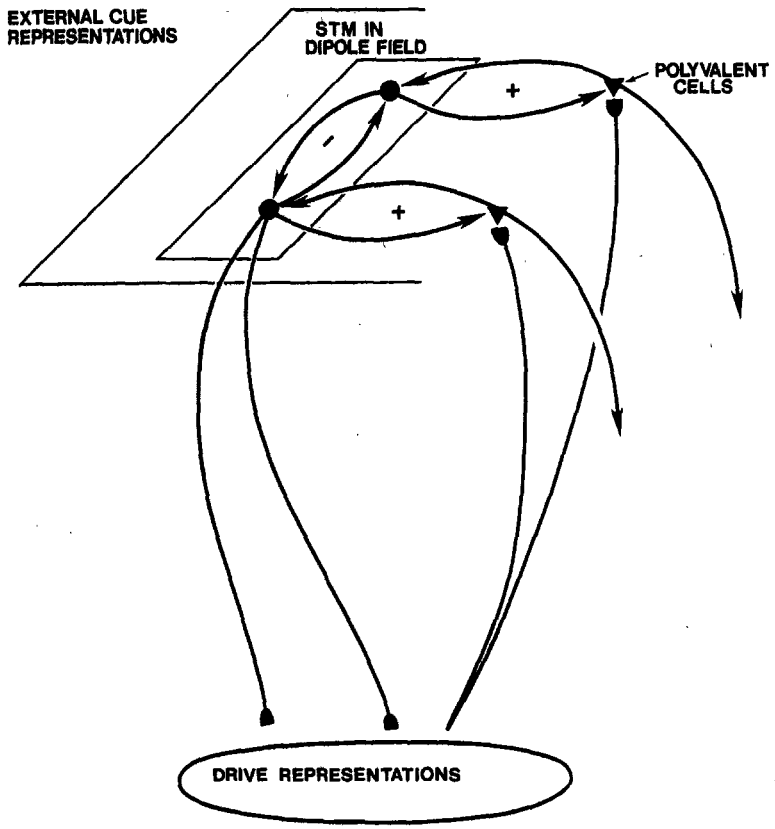

Figure 13. Interaction of external cue and incentive motivation signals at polyvalent cells. Let a set of nodes, or cells, in the dipole field be activated by an external scene. A pattern of short-term memory activity across the nodes represents the scene. Each such node sends an excitatory signal to its polyvalent node, or cell. Signal size is an increasing function of short-term memory activity. These specific signals are insufficient to fire the polyvalent cells. Sufficiently large incentive motivational signals from a drive representation must simultaneously converge on 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 thereby fire a subset of the polyvalent cells that are activated by the external scene. The relative rate of firing of each polyvalent cell will depend jointly on the short-term memory activity of its trigger cues in the scene and on the relative size of its long-term memory trace in the conditioned reinforcer pathway. When a polyvalent cell fires, it delivers positive feedback signals to the cue-activated cells that supply it with specific short-term memory signals. This positive feedback from polyvalent cells selectively augments the short-term memory activities of certain cue-activated cells, which thereupon can more strongly inhibit the short-term memory of other representations in the dipole field using the short-term memory normalization property. The incentive motivational properties of certain cues thereby alter the set of cues to which the network pays attention. The polyvalent cells that can maintain their firing can also read out learned patterns (e.g., motor commands) to other parts of the network.

for my later work on reinforcement and motivation.

For present purposes, I need to emphasize one difference between Figure 12 and Figure 13. The anatomy in Figure 13 separates the firing of polyvalent cells from the STM reverberation through gated dipoles. Owing to this property, a sensory representation can reverberate in STM and thereby deliver signals to a polyvalent cell, or cells, without firing these cells. A polyvalent cell in Figure 13 can fire only if it simultaneously receives STM signals from an external cue representation and incentive motivation signals from a drive representation. This property is analogous to John's (1966, 1967) reports that

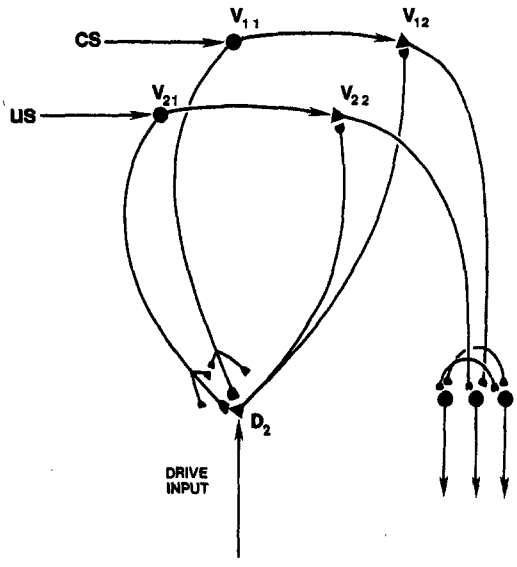

Figure 14. Path equivalence of CS and US representations. Both the conditioned stimulus (CS) and the unconditioned stimulus (US) activate similar network designs. This path equivalence property overcomes the asymmetry of CS and US contributions to the modified Rescorla-Wagner Equation 8. Firing of the polyvalent cells V_{12} and V_{22} is prevented except when sufficiently large specific signals from V_{11} and V_{21}, respectively, and nonspecific signals from D_2 simultaneously converge at the polyvalent cells. The network stages $\{V_{11}, V_{12}\}$ are part of the CS representation. The stages $\{V_{21}, V_{22}\}$ are part of the US representation. The CS and US activate similar network anatomies, but the LTM traces in these anatomies are not identically distributed. In particular, a US-activated signal from V_{21} to D_2 can fire D_2 if the drive input is sufficiently large. A CS-activated signal from V_{11} to D_2 cannot fire D_2. See the text for how CS–US pairing endows the CS with conditioned reinforcer, incentive motivation, and habit readout capabilities.

certain polyvalent cortical cells that are involved in cortical conditioning can fire only in response to a sum of CS and US signals. The property is also analogous to the effects of anodal DC potential shifts on cortical conditioning (Morrell, 1961; Rusinov, Note 1). In my theory, the anodal DC shift replaces the requirement of an incentive motivation signal to fire polyvalent cortical output cells.

33. The Multidimensional Nature of Secondary Conditioning

The functional separation of STM reverberation and polyvalent cell firing implies the following description of how a CS acquires US properties to control observable behavior (Grossberg, 1971). This description shows how to overcome the asymmetry between CS and US in the modified Rescorla–Wagner Equation 8, explains the Dickinson and Mackintosh (1979) data on selective effects of distinct reinforcers on associability, and shows that drive representations are functionally separate "processors" from cue representations, in contrast with the Pearce and Hall (1980) theory.

Let a CS activate the population V_{11} (Figure 14), which thereupon begins to reverberate in STM. Then V_{11} sends specific signals to the polyvalent cell population V_{12} (among others) and nonspecific signals to the drive representations. Nothing else happens until a US arrives at population V_{21}. This is because V_{12} can fire only if it receives an input from V_{11} and an incentive motivation input from a drive representation, but the signal from V_{11} to the drive representations is initially too small to fire them. When the US perturbs V_{21}, V_{21} sends signals to the polyvalent cells V_{22} and to the drive representations. These latter signals can fire a certain drive representation, if its drive input is large enough, because the cue firing V_{21} is a US for that drive representation, which I will henceforth denote by D_2. When D_2 fires, it releases nonspecific incentive motivation signals to all the external cue representations. Now five things happen.

First, because V_{11} and D_2 are both active, the LTM traces in the pathway from V_{11} to D_2 are strengthened. When these LTM traces get strong enough, the CS alone will be able to fire D_2. Second, the nonspecific incentive motivational signal from D_2 combines with the US-derived signal from V_{21} at V_{22}, thereby firing polyvalent cell signals from V_{22}, which read out an unconditioned response (UR) pattern. Third, because the incentive motivation signal is nonspecific, it also combines with the CS-derived signal from V_{11} at V_{12}, thereby firing the polyvalent cells V_{12}. Fourth, because D_2 and V_{12} are both active, the LTM traces in the pathway from D_2 to V_{12} are strengthened. Fifth, the polyvalent cells V_{12} fire sampling signals to the cells at which the UR pattern is being read out. These signals encode (a fractional component of) the UR in the LTM traces of this pathway. The en-

coded pattern will henceforth be read out as a conditioned response (CR) pattern. The CS thereby acquires US properties owing to LTM encoding in conditioned reinforcer, incentive motivation, and habit strength pathways.

This network provides a simple answer to the synchronization question: How does the US turn on the CS with just the right time lag to sample read out of the UR pattern? The same incentive motivation burst that allows V_{22} to read out the UR also allows V_{12} to emit sampling signals that read in the CR.

In the remaining sections I will use the properties of adaptive resonance and reset in the cognitive circuits (Section 19) and in the cognitive-motivational circuits (Sections 30–33) to suggest explanations of conditioning and attentional data, including data that no single formal model can explain.

34. Unblocking, Context, and Habituation

Unblocking is produced by the surprising omission of a second shock (Dickinson, Hall, & Mackintosh, 1976). In my theory this is because the active internal representations of previously presented cues read out a learned expectancy whose mismatch by the unexpected event triggers dipole reset and consequent STM enhancement of previously overshadowed cues (Section 23).

Experiments might be designed to test how the amount of rebound or enhancement depends on the number of cues that have been stored in STM when the unexpected event occurs. Due to the limited-capacity restriction imposed by STM normalization, storing more cues can reduce the STM activity of each cue and thereby reduce the rate of transmitter depletion in each dipole gate. The rebound pattern in response to a fixed arousal burst will therefore change as a function of the number of active representations. The simplest version of this idea suggests that if more cues are simultaneously stored in STM for a given amount of time, then they will each be rebounded less by an unexpected event. Let us suppose that a less rebounded cue representation retains more STM activity than a more rebounded cue representation and that this residual activity can summate with the STM activity elicited by a later presentation of the same cue. Let two experimental groups differ according to how many cues are stored in STM and expose both groups to an unexpected stimulus array that includes one of the previously stored cues and a new cue. Suppose that the unexpected stimulus array resets STM by rebounding or enhancing the previously stored representations. Other things being equal, a less rebounded cue may preserve more residual STM activity for summation with its reoccurrence as part of the unexpected event. If this happens, the cue's total STM activity will be larger if it was part of a larger set of previously stored items than a smaller set of items. Hence, it will be better attended in the presence of the new cue.

However, other things are not usually equal. Storing more cues simultaneously may provide each cue with less STM activity due to the normalization property. Hence, the occurrence of a smaller reset per cue may provide no advantage to larger sets of cues. By contrast, storing more cues might cause more of their STM activities to be enhanced, rather than rebounded, by an unexpected event. When this occurs an advantage may indeed accrue to larger sets of cues. Both of these effects will be sensitive to the duration with which the cues are stored before they are reset. Large STM activities deplete their transmitter gates faster than small STM activities. Consequently, the relative disadvantage to smaller sets of stored cues may be greater if the storage duration is shorter. Finally, a switch from more stored cues to a new event that includes only one of these cues may cause a greater mismatch, and hence be more unexpected, than a switch from fewer stored cues to the same new event, although normalization tends to counter this effect also. The reset arousal burst that occurs after more cues are stored may thus be larger and may offset any advantage due to slower transmitter habituation. Parametric studies are needed in which the number of cues originally stored, the duration of storage, the number of new cues to be stored, and the amount of overlap between the two sets of cues are varied to disentangle the relative contributions of STM normalization, mismatch, and reset mechanisms on the reallocation of attention.

35. Double Shock Experiments: Initial STM Advantage of Surprising Events

To explain the interesting experiment of Mackintosh, Bygrave, and Picton (1977) (Section 5), I accept the fact that on the first compound trial, the tone is at its worst as a predictor of shock, having never before been paired with shock. I hereby avoid the internal problem within Mackintosh's theory (Section 3). Because the tone has never before been presented, it is most unexpected on the first compound trial. More precisely, on the first trial the tone possesses its maximal potency for mismatching the learned sensory expectancies that are operative in the situation, in particular, the expectancies that are controlled by situational cues. Consequently, tone onset triggers a relatively large STM reset that abets the STM storage of the tone's sensory representation. This is the first main point: The tone's very unexpectedness can enhance its initial STM storage without regard to what its STM representation samples later on. On successive tone presentations, this intitial advantage will fade, other things being equal, as the tone is incorporated into the pattern of learned feedback expectancies.

After the tone's advantageous initial STM storage takes place, the tone's STM representation begins to emit several types of signals. Some of these signals initiate the process whereby the tone is incorporated into higher order chunks (Section 18). Others of these signals begin to sample the drive representations (Section 30). On the first compound trial, the light can also send negative conditioned reinforcer signals to the drive representation with which it was previously associated.

An important issue is: How much are the light's conditioned reinforcer signals reduced by the occurrence of the tone? For example, such a reduction can occur because of a direct STM competition between light and tone representations via the STM normalization effect. An indirect reduction can be due to antagonistic rebound of the active chunks that bind situational cues together into context-sensitive representations. Such a rebound can be triggered by the arousal burst that is contingent on the tone's unexpectedness in a given experimental context. Inhibiting these chunks can eliminate the subliminal feedback signaling to expected situational cues, including the light, and can therefore decrease the light's STM activity by removing a source of STM matching. This issue also arises on Trial 25 in Group Y of the Kamin (1969) experiment reviewed below (Section 45).

36. Conditioned Reinforcer Learning

After both the tone STM representation and the light STM representation send signals to the drive representations, the light representation can supraliminally activate the negative drive representation that was previously activated by shock on conditioning trials, because the light is now a conditioned reinforcer (Sections 27 and 33). Once this drive representation is supraliminally activated, the tone's sampling signals can acquire some negative conditioned reinforcer properties via LTM encoding during the first moments after tone–light onset. This LTM change can occur whenever the tone's sampling signals are contiguous with supraliminal activity at the drive representation. On the first compound trial, the tone representation emits unusually large sampling signals because it has acquired unusually large STM activity due to its unexpectedness. Large sampling signals cause their LTM traces to encode sampled STM patterns faster than do small sampling signals. The tone can therefore acquire negative incentive properties much more rapidly on a trial when it is unexpected than on a trial when it is expected. This fact explains why the tone can condition so well—as a conditioned reinforcer—on the trial when it is first presented.

37. Incentive Motivation Feedback Influences STM Competition

As the tone begins to acquire negative reinforcing properties by being conditioned to the negative drive representation, the firing of this drive representation also releases conditioned incentive motivation signals preferentially to the light representation (Figure 13). More precisely, the light representation activates its polyvalent cells directly and via the conditioned reinforcer-incentive moti-

vation loop through the drive representations. Then these polyvalent cells can fire. The polyvalent cells thereupon feed excitatory signals back to the light representation to further enhance its STM activity. Because of STM normalization among the cue representations, the enhanced STM of the light representation competitively inhibits the STM of the tone representation. Whether this feedback inhibition can entirely shut off the tone representation depends on how many and how intense the prior light–shock trials were (and thus how large and selective the conditioned incentive feedback signals are) and on how surprising the tone was (and thus how big an advantage it acquired in its initial STM storage).

If the tone representation is entirely suppressed, it will have acquired negative conditioned reinforcer properties but no conditioned incentive-motivation feedback to its own sensory representation. Even this eventuality should not prevent the tone from eliciting fearful reactions when it is later presented alone. Because the light is not then present to suppress it, the STM normalization property allows the tone to be stored in STM, whereupon the tone's conditioned reinforcer signals can elicit a fear reaction.

38. Salience Influences Overshadowing

This explanation also shows why, when the two CSs differ markedly in their intensity or salience, the more salient cue can overshadow the less salient cue (Mackintosh, 1976). In my theory the greater cue saliency or intensity gives it an initially larger STM strength, larger LTM sampling signals between cue and drive representations, and thus a competitive advantage from the very start.

39. Overshadowing on a Single Trial

These concepts suggest how overshadowing effects can occur when only a single trial of compound conditioning is given (Mackintosh, 1971; Mackintosh & Reese, 1979). As seen above, the tone's best chance to avoid overshadowing is to achieve strong initial STM storage. Without an initial advantage the tone's conditioned reinforcer learning will be slow at best in the time interval before the light's incentive feedback enables it to suppress the tone's STM. I have also indicated above several ways to abet tone conditioning.

We are now ready to consider why a second shock 10 seconds after a first shock on the first compound trial need not make the tone more fearful. Also, we will see why two shocks on one trial followed by either one or two shocks on a second trial can lead to a more fearful tone than one shock on the first trial followed by either one or two shocks on a second trial.

40. The Tone Is More Unexpected After Two Shocks

If the tone is not on when the second shock occurs, then the tone's STM representation may not send sampling signals to the drive representation when it is activated by the second shock. Thus, the tone does not acquire more negative conditioned reinforcer strength because of the second shock on the first compound trial. Why then does the tone acquire significantly more negative conditioned reinforcer strength on the second compound trial? This can be explained by noting that the second shock on the first compound trial occurs after a series of trials during which only one shock occurs. Thus, the second shock is unexpected. Also, the second shock occurs after the tone has unexpectedly occurred. The tone's unexpected occurrence initiates the process whereby the tone representation is incorporated into the pattern of situational expectancies. The occurrence of the second shock then alters the pattern of learned situational expectancies beyond the alterations already triggered by the tone. Consequently, when a tone occurs on a second compound trial that follows two shocks, it is more unexpected than a tone that occurs on a second compound trial that follows one shock. Due to the tone's greater unexpectedness, the tone's initial STM storage on the second compound trial is greater, its sampling signals to the drive representations are larger, and its rate of conditioned reinforcer learning is accelerated. An independent test of the tone's greater unexpectedness would be achieved if the tone elicits a larger P300

evoked potential when it occurs after two shocks than after one shock.

This description of how a tone achieves its superior STM storage after two shocks than after one shock suggests that the light may also be better stored in STM after two shocks than after one shock, because both cues disconfirm the second shock component of the situational expectancies. A greater initial STM storage of the light does not prevent the tone from strengthening its negative reinforcer strength on the second compound trial for two reasons: (a) The relative advantage of the light is not greater than that of the tone, so the light will not competitively inhibit the tone via the normalization property during the phase of initial STM storage. (b) Once both the light and tone representations begin to send signals to the drive representations, the larger signals emitted by the light representation can speed up the tone's conditioned reinforcer learning by increasing the activity of the negative drive representation. One way to test the effect of the second shock on the initial STM storage of the light is to measure the P300 evoked potential that is triggered if the light alone, rather than a light-tone compound, is presented after one or two shocks.

41. Situational Cues

Having come this far, we are now ready to raise an issue that Pearce and Hall (1980) do not mention. When the second shock occurs on the first compound trial, it is a surprising event whose negative reinforcing properties will be conditioned to simultaneously active cue representations. These representations will include the representations of situational cues that are again present when the tone is presented on the test trial. Why does the tone-plus-situational cue readout of negative conditioned reinforcer signals not create more negative incentive after the second shock than it does after the first shock? My answer is that the surprising occurrence of the tone on the test trial tends to suppress the STM of the situational cues via antagonistic rebounds. Then the tone's STM will tend to control the net conditioned reinforcer readout from attended sensory representations. To test this explanation, one might try covarying two experimental properties: the intensity of the second shock and how surprising the tone is on the next test trial.

42. Release From Overshadowing by an Unexpected US

In experiments wherein the tone is repeated during several compound trials, the tone's initial STM advantage due to its unexpectedness wears off as it is incorporated into the pattern of situational expectancies. The tone's conditionability thereby fades. If an unexpected US follows a light–tone combination, the tone's STM activity can be amplifed owing to the differential enhancement of overshadowed representations by the STM reset event (Section 23). The US also activates a drive representation. Because of the simultaneity of enhanced STM activity in the tone representation and in the drive representation, the tone can acquire both conditioned reinforcer and incentive motivation properties. These LTM changes are not rapidly terminated by competitive signaling from the light representation because the STM of this representation has been attenuated by the reset event.

If the unexpected US reoccurs on several trials, its unexpectedness also fades. By the time this happens, however, the unexpected US has endowed the tone's representation with a conditioned positive feedback loop to itself via the drive representation. A shift gradually occurs as trials proceed from the tone's initial STM advantage due to the shock's unexpectedness—which is mediated by situational expectancies and the orienting subsystem—to a more lasting LTM advantage due to the tone's reinforcing and incentive motivation properties—which manifest themselves as an attentional resonance.

43. Modulation of US and Drive Input Effects by Expectancies, Conditioned Reinforcers, and Transmitter Habituation

A further remark needs to be made about which drive representation is activated by the shock US. This is a subtle matter because it depends on US intensity, the degree of US expectedness, and the conditioned reinforc-

ing potency of the light representation due to prior learning.

Were a shock to suddenly turn on out of context, it would certainly activate a negative drive representation (Section 28). This need not happen if a shock turns on while the light representation is on. This is true because the light representation is already sending conditioned reinforcer signals to the negative drive representation when the shock occurs. The unexpectedness of the shock attenuates the STM activity of the light representation. A sudden reduction in conditioned reinforcing signals to the negative drive representation is thereby caused. The shock can offset this reduction in input by generating its own unconditional input to the negative drive representation. If the shock input is larger than the prior level of conditioned reinforcing signals, then the total input to the negative drive representation will increase and a fear reaction will be elicited. If, however, the shock-induced input is smaller than the conditioned reinforcer decrement, then the total input to the negative drive representation will suddenly decrease. The shock can thereby cause an antagonistic rebound that activates the positive drive representation that shares a gated dipole with the negative drive representation (Section 13). The onset of shock can thereby cause a transient relief reaction. This argument also indicates how an unexpected increase in a shock US can cause the tone to become a negative reinforcer, whereas an unexpected decrease in the shock US can cause the tone to become a positive reinforcer (Kamin, 1968, 1969; Rescorla, 1969; Wagner, 1969; Weisman & Litner, 1969), despite the fact that the shock US activates a negative drive representation in both cases. Given a fixed decrease in shock intensity, the rebound size should be an increasing function of shock unexpectedness (measured perhaps as a larger P300 evoked potential) and an increasing function of the conditioned reinforcer strength of the light (measured perhaps by the number of preceding light–shock trials). The emotional meaning of the shock is thus determined by an interaction of its unconditional input with the pattern of active expectancies and conditioned reinforcing signals at the moment of its occurrence. In Grossberg (1982a) I propose that a similar argument, wherein hunger drive input replaces shock input, explains some paradoxical findings about eating and satiety. Oropharyngeal signals that are gated by conditioned reinforcer LTM traces are suggested to alter the effects of hunger drive input much as light-induced signals that are gated by conditioned reinforcer LTM traces are suggested to alter the effect of shock input.

Varying the suddenness with which a shock is turned on can alter these conclusions by influencing both the expectancy and the reinforcing properties of shock. One effect of shock onset rate on reinforcement is the following: Suppose that a shock slowly turns on from an initial intensity of 0 to a final intensity of J. Because the shock increase is gradual, the transmitter in the on-channel of the gated dipole is gradually depleted, or habituates, at a rate proportional to signal strength times the amount of available transmitter (Section 13). Because the transmitter level accumulates slowly, by the time the shock intensity J is reached, the amount of transmitter Z can be much smaller than its maximal amount B. Thus the effect of intensity J is proportional to $f(J + I)Z$, where I is the arousal level, and $f(w)$ is the sigmoid signal (Section 22). By contrast, a sudden shock creates the signal $f(J + I)B$, where $B > Z$, because transmitter is fully accumulated when the shock intensity suddenly switches from intensity 0 to intensity J. Sudden shocks can thereby be more negatively reinforcing than gradually increasing shocks (Church, 1969, Miller, 1960). In Grossberg (1982a) I suggest a similar argument about transmitter habituation rates in gated dipoles to explain drug tolerance and withdrawal effects, including symptoms like rebound insomnia. Many of the expectancy, reinforcing, and transmitter habituation effects that occur during conditioning experiments have mechanistic analogs in other behavioral and clinical syndromes.

44. Latent Inhibition

Similar concepts can be used to explain the following interesting Hall and Pearce (1979) experiment. In Stage 1 of this experiment, a tone was paired with a weak shock

in the experimental group. In the control group a light was paired with the same shock. In the next stage the tone preceded a stronger shock in both groups. In the experimental group, learning was slower. In my theory this occurs because the tone is more unexpected in the control group, thereby acquiring a greater initial advantage in STM, and therefore conditions faster, as in Section 43.

Rather than continue to explain other data that Pearce and Hall (1980) mention, I will suggest an explanation of some classical experiments that seem to go beyond the capabilities of all the formal theories. I will also suggest an interdisciplinary paradigm to test my explanations.

45. Interaction of STM Reset, STM Normalization, and LTM Gating in CER Experiments

The STM normalization property often holds only partially, due to the fact that feedback inhibitory interactions can decrease as a function of intercellular distances. These distance-dependent interactions help to define the generalization gradients that determine which cues mutually interact during conditioning experiments (Grossberg, 1975, 1981b). A possible instance of a partial STM normalization effect due to Pavlovian conditioning is the somewhat faster learning of a conditioned emotional response to a compound stimulus than to its component stimuli (Kamin, 1969). Parametric studies of compound training trials using approximately equally salient stimuli whose similarity is parametrically altered across experiments, followed by extinction of the compound, or of each component taken separately across groups in each experiment, would provide useful theoretical information about the interaction between the degree of STM normalization and the rate of CER learning.

Another piece of data reported in Kamin (1969) also suggests an STM normalization effect. His Group Y first learned a CER to a noise stimulus (N) on 16 trials, then received a nonreinforced compound light–noise (LN) stimulus for 8 trials, and finally received 4 more nonreinforced N trials. His Group Z also received 16 CER trials with stimulus N, but these trials were follows by 12 nonreinforced trials with N. Three main effects were found: The first LN trial in Group Y showed a larger suppression ratio than the first nonreinforced N trial in Group Z. The suppression ratio increased on successive nonreinforced trials of LN in Group Y and of N in Group Z. On the first nonreinforced N trial in Group Y (its 25th trial), the suppression ratio suddenly dropped to the value that it had on the first nonreinforced LN trial. This suppression ratio was, moreover, significantly lower than the suppression ratio on Trial 25 in Group Z.

Kamin was impressed by the rapidity with which the suppression ratio changed on the first nonreinforced LN trial and on the first nonreinforced N trial for Group Y. He realized that the Y animals rapidly noticed L and that their processing of L somehow attenuated the suppression. In my theory the surprising occurrence of L abets its STM storage, weakens the STM storage of N via STM normalization, and thereby reduces the negative conditioned reinforcing signals from L to the drive representations.

That an STM rather than an LTM effect is primary on the transitional trials is further suggested by what happens on Trial 25 in Group Y. When N is then presented without L, its representation can acquire a larger STM activity. This representation can then read out—on that very trial—a larger negative conditioned reinforcing signal to the drive representations. The negative reinforcing LTM trace is there to be read out because the extinction of the N representation on LN trials was slowed owing to its small STM activity.

46. Overshadowing During Key Pecking

Newman and Baron (1965) reinforced pigeons who pecked a vertical white line on a green key (the $S+$) but not a green key alone (the $S-$). They tested cue discrimination by tilting the line at various orientations during recall trials. A generalization gradient of pecking was found, indicating that the vertical line was discriminated. By contrast, no generalization gradient was found if the $S-$ on learning trials was a red key or if the $S-$ was a vertical white line on a red key.

Newman and Benefeld (cited in Honig, 1970) used a vertical white line on a green key as $S+$ and a green key as $S-$ but tested and found generalization of the line orientation on a black key. They also tested generalization on a black key following training without a green $S-$ and again found a generalization gradient, by contrast with the case where testing used a green key. They interpreted this effect as "cue utilization during testing rather than cue selection during learning" (p. 202). This interpretation does not explain how the orientation cue could be learned on training trials if it was not discriminated using a green background on test trials yet could be discriminated using a black background on test trials if it was not learned on training trials.

My explanation of these data begins by noting that color cues are prepotent over orientation cues in the pigeon, other things being equal. Consequently, when a vertical white line on a green background is first presented, the green representations will partially overshadow the orientation representations. (I will talk about "green" and "orientation" representations as a shorthand for more sophisticated coding notions that we do not need here.) Grossberg and Levine (1975) and Levine and Grossberg (1976) describe some factors that control how prepotent representations can mask the STM activities of other representations due to competitive feedback interactions.

When the line-on-green cues are first presented, they enjoy an additional advantage in their STM storage. Their unexpectedness in the context of the experiment's situational cues will strengthen the STM activities of the line-on-green cues as the STM activites of the situational cue representations are rebounded. These rebounds should elicit a P300 evoked potential.

After the line-on-green representations are initially stored in STM, the green cues can increase their relative STM advantage as they acquire conditioned reinforcer and conditioned incentive motivation properties. They do this by means of the conditioned-reinforcer–incentive-motivation loop, the polyvalent cell firing constraint, and the STM normalization property in the manner described within Section 37.

The orientation representations can also acquire conditioned reinforcer and incentive motivation properties just so long as their STM activities are not suppressed. Their learning rates will be slower than those of the green representations, because their sampling signals in the conditionable pathways are smaller due to their smaller STM activities. Hence, their conditioned pathways will remain weak compared to those of the green representations. As conditioning continues, the orientation representations may be entirely suppressed if the conditioned advantage of the color cues becomes sufficiently great to drive orientational STM activities to zero by competitive feedback across the cue representations.

The unexpected nonoccurrence of reward in response to pecking the green key causes an antagonistic rebound that excites the off-cells of the previously most active STM representations. The active incentive motivation pathways thereupon sample a large off-response in the green representational dipoles (Figure 11). As this experimental contingency reoccurs on several trials, the *net* incentive motivation feedback to the green dipoles becomes progressively smaller due to dipole competition between the conditioned on-cell and off-cell pathways to these dipoles. This is just the extinction mechanism of Section 28 acting at the sensory representations rather than at the drive representations.

Even zero net incentive feedback may not be small enough to extinguish the green representation, however, because of the innate advantage of color over orientation. Negative net incentive feedback may be needed to overcome green's innate competitive advantage. Net negative feedback is needed if net positive conditioned reinforcer-incentive feedback to the orientation representation is not sufficient to offset the innate competitive advantage of the color representation when the latter receives net zero conditioned feedback.

This framework explains why the white vertical line is discriminable on a black background during test trials even if it is not discriminable on a green background during test trials in an experiment without a green $S-$ on learning trials. Removing green on test trials eliminates competitive feedback from

the color representations to the orientation representations. The STM field is thereby renormalized. In the renormalized field, even small conditioned-reinforcer–incentive-motivation feedback signals can provide the white vertical line representation with a significant competitive advantage for STM storage.

47. The Problem of Perseverating Prepotent Cues

The above discussion shows how the conditioned - reinforcer – incentive - motivation feedback loop enables representations to overcome innate competitive STM disadvantages. Some further remarks might clarify why the incentive motivation pathway must send branches to both the on-cells and off-cells of cortical dipoles, just as the conditioned reinforcer pathway sends branches to both the on-cells and off-cells of the drive representation dipoles. The main benefit is that some cues can lose net positive feedback as other cues gain net positive feedback while both sets of cues are conditioned to the same drive representation. This property avoids the following dilemma:

Suppose the rebound that conditions zero net feedback to the green representation occurs among the drive representations rather than among the cue representations. Then rebound activates a negative drive representation, and the net conditioned reinforcer output controlled by the green representation becomes small, rather than the net incentive motivational output driven by a large conditioned reinforcer output becoming small, as in Section 46. This mechanism is unstable for the following reason: As soon as the orientation representation takes over in STM, its positive conditioned reinforcer signals activate the positive drive representation. When this drive representation sends incentive motivation feedback to the cortex, the green representation receives conditioned positive feedback because the negative drive representation is momentarily inhibited. Then the green representation can quickly overshadow the orientation representation because of its innate competitive advantage. As soon as the green representation is reinstated in STM, its conditioned reinforcer signals cause readout of net negative incentive from the competing drive representation. The green representation is consequently shut off, the orientation representation is disinhibited, and the cycle repeats itself.

Any viable alternative to the present network description must also avoid this problem of perseverating prepotent representations. In particular, a more sophisticated coding analysis would replace "green" representations and "orientation" representations with heterarchical network encodings wherein one representation's prepotence for masking other representations would depend on its heterarchical position with respect to all representations. In Grossberg (1978b, Sections 25–47) I suggest some rules whereby heterarchical masking can be designed to choose those chunks that provide the most informative prediction in a prescribed cue context during recognition and recall.

48. Two Distinct P300 Evoked Potentials in Cortex and Hippocampus

The previous discussion of overshadowing during key pecking suggests a striking psychophysiological prediction. I have argued that the unexpected nonoccurrence of reward in response to pecking the green key can gradually extinguish the net incentive motivation to the green representation. This occurs as the incentive motivation LTM traces sample the antagonistic rebound within the green representation on successive nonreinforced trials. The reset of the green representation also has an immediate effect on each trial. Offset of this representation rapidly shuts off conditioned reinforcer input to the positive drive representation. If the green representation has been conditioned to the drive representation on sufficiently many previous reinforced trials, the reduction in conditioned reinforcer input will be large enough to overcome the STM hysteresis that defends the positive drive representation against reset. Then a rebound will occur within the drive dipole itself, thereby activating its negative drive representation. If a new cue is stored in STM at the time of this negative drive rebound, it will become a negative conditioned reinforcer by sampling the rebound.

Let us consider the minimal assumption that any massive antagonistic rebound in a

catecholamine dipole system is registered as a P300 evoked potential—keeping in mind that rebounds in different brain regions may occur yet ultimately be associated with distinct evoked potentials. Then the above discussion predicts that the nonoccurrence of expected reward can trigger a cortical P300 by mismatching a learned expectancy. The reset of cortical STM can thereupon trigger a hippocampal P300 by rapidly withdrawing conditioned reinforcer input. The size of this second P300 should, moreover, depend on the strength of the conditioned reinforcer due to the number of preceding conditioning trials. If these P300 predictions hold up, they will clarify that a P300 can be elicited by different operations in different brain regions. They will also refine our understanding of the information processing substrates of overshadowing and discrimination learning by distinguishing rebounds that motivationally extinguish cues owing to their cognitive irrelevance from rebounds that directly elicit new conditioned reinforcer learning.

49. Nonmonotonic Relation Between P300 and CNV

The key pecking experiment also suggests a psychophysiological test that would argue for P300 as a measure of STM reset and against P300 as a measure of LTM learning. I suggest that the unexpected occurrence of a reward in response to the line-on-green cue will elicit a P300. As this P300 shrinks on successive rewarded trials, I suggest that the line-on-green cue will elicit a growing motivational CNV that reflects the progressive conditioning of positive net incentive motivation (Cant & Bickford, 1967; Irwin, Rebert, McAdam, & Knott, 1966). By contrast, I suggest that the unexpected nonoccurrence of a reward in response to green alone will elicit a P300. As this P300 shrinks on successive unrewarded trials, the green cue should elicit a shrinking motivational CNV as the net incentive motivation of the irrelevant green cue is extinguished. In the former case a conditioned response is learned, whereas in the latter case a conditioned response is extinguished. If these predictions are verified, then we can conclude that P300 size does not differentiate opposite outcomes in LTM because a monotonic decrease in P300 can predict either CNV increase (learning) or CNV decrease (extinction) within the same experiment.

50. Some Comparisons With the Formal Models

Pearce and Hall (1980) ascribe extinction to competition between CS–US and CS–$\overline{\text{US}}$ ($\overline{\text{US}}$ = no US) associations due to an "inhibitory link between the US and $\overline{\text{US}}$ memories" (p. 546). They suggest this concept to replace Rescorla and Wagner's (1972) notion that extinction is due to weakening of previously established associations. My own concept of how a conditioned reinforcing cue's input to a gated dipole is extinguished is superficially similar to Pearce and Hall's (Grossberg, 1972a, 1972b). I also suggest, however, that the competitive extinction process is mediated by the drive representations and is due to gated dipole rebounds. Neither these concepts nor their mechanistic substrates appear in the formal models.

Instead, the formal models restrict themselves to links between CS and US memories, which in turn read out the CR. In my theory, readout of the CR does not require activation of a US memory, but only of the LTM-encoded patterns that were sampled by the CS from STM when the US was active, as in Sections 15 and 18. These LTM-encoded patterns can be a fractional component or other transformation of the US, due to nonisotropy of CS and US sampling pathways across the network or due to STM transformations of the US pattern before it is encoded in LTM at the CS-activated synaptic knobs. I do not see how direct links from CS to US can account for the sometimes significant differences between UR and CR, whereas an STM-mediated theory can easily do so (Seligman & Hager, 1972).

Pearce and Hall (1980) suggest "that a $\overline{\text{US}}$ representation is activated only by the omission of an *expected* US" (p. 543) and suggest a formula for the intensity $\bar{\lambda}$ of the reinforcer $\overline{\text{US}}$, namely,

$$\bar{\lambda} = V_\Sigma - \bar{V}_\Sigma - \lambda. \tag{14}$$

I agree that an off-cell rebound can be acti-

vated by the nonoccurrence of an expected event, mediated by a mismatch-contingent arousal burst. However, this is not the only way to activate a US in my theory. Just as sudden offset of shock can trigger relief (Denny, 1971), a rebound can also be caused by the mere offset of a reinforcer. Furthermore, the Equation 14 for rebound size is inadequate for many reasons: It does not explain the inverted U in reinforcement (Berlyne, 1969). It does not explain the analgesic effect whereby cutting J units of shock in half is less rewarding then shutting $J/2$ units of shock off (Campbell & Kraeling, 1953). It does not explain why the reinforcing effect of shock offset should depend on the prior duration of shock (Boe, 1966; Borozci, Storms, & Broen, 1964; Church, Raymond, & Beauchamp, 1967; Keehn, 1963; Strouthes, 1965). All of these properties obtain in gated dipoles (Grossberg, 1972b). Moreover, Equation 14 lumps together LTM expectancy, STM matching, nonspecific arousal, and STM rebound properties in a way that obscures their mechanistic substrates, notably their influence on STM and LTM patterns rather than parameters.

At bottom, the formal models are led to these difficulties because they do not adequately distinguish the STM and LTM mechanisms that are used in the processing of expected and unexpected events. Consequently, the formal models cannot easily make the distinction that a surprising CS can reset STM in a manner that favors its own subsequent STM storage, whereas a fully predictable US can also be stored (or, as Pearce and Hall would say, "processed") in STM by resonating with an active feedback expectancy. The recent theorizing of Wagner (1978) on STM priming perhaps comes closest to making these distinctions within the stream of formal models.

51. Schedule Interactions and Behavioral Contrast

Similar difficulties occur in recent models of instrumental conditioning. Instead of overemphasizing LTM properties at the expense of STM properties, the Hinson and Staddon (1978) theory of schedule interactions and behavioral contrast completely forsakes LTM effects to focus on STM competitive properties. In Grossberg (1981b) I show that the same theoretical ideas that I sketched herein can overcome some difficulties that Hinson and Staddon face in explaining their data, and I make some predictions to test these ideas.

Pavlovian and instrumental experiments that have heretofore been analyzed in a fragmentary fashion by formal models, at the cost of implying internal paradoxes and restricting their predictive power, can be understood in a unified fashion in terms of a few psychophysiological mechanisms whose existence can be more directly validated by interdisciplinary experimental paradigms.

Reference Note

1. Rusinov, V. S. *An electrophysiological analysis of the connecting function in the cerebral cortex in the presence of a dominant area.* Paper presented at the XIX International Physiology Congress, Montreal, 1953.

References

Anand, B. K., & Pillai, R. V. Activation of single neurones in the hypothalamic feeding centres: Effect of gastric distension. *Journal of Physiology*, 1967, *192*, 63–77.

Beach, F. A. Characteristics of masculine "sex drive." In M. R. Jones (Ed.), *Nebraska Symposium on Motivation* (Vol. 4). Lincoln: University of Nebraska Press, 1956.

Berlyne, D. E. The reward-value of indifferent stimulation. In J. T. Tapp (Ed.), *Reinforcement and behavior.* New York: Academic Press, 1969.

Boe, E. E. Effect of punishment duration and intensity on the extinction of an instrumental response. *Journal of Experimental Psychology*, 1966, *72*, 125–131.

Bolles, R. C. *Theory of motivation.* New York: Harper & Row, 1967.

Borozci, G., Storms, L. H., & Broen, W. E. Response suppression and recovery of responding at different deprivation levels as functions of intensity and duration of punishment. *Journal of Comparative and Physiological Psychology*, 1964, *58*, 456–459.

Brown, J. L. Afterimages. In C. H. Graham (Ed.), *Vision and visual perception.* New York: Wiley, 1965.

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

Campbell, B. A., & Kraeling, D. Response strength as a function of drive level and amount of drive reduction. *Journal of Experimental Psychology*, 1953, *45*, 97–101.

Cant, B. R., & Bickford, R. G. The effect of motivation on the contingent negative variation (CNV). *Electro-*

encephalography and Clinical Neurophysiology, 1967, *23*, 594.

Carpenter, G. A., & Grossberg, S. Adaptation and transmitter gating in vertebrate photoreceptors. *Journal of Theoretical Neurobiology*, 1981, *1*, 1–42.

Church, R. M. Response suppression. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Church, R. M., Raymond, G. A., & Beauchamp, R. D. Response suppression as a function of intensity and duration of punishment. *Journal of Comparative and Physiological Psychology*, 1967, *63*, 39–44.

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

Dexter, W. R., & Merrill, H. K. Role of contextual discrimination in fear conditioning. *Journal of Comparative and Physiological Psychology*, 1969, *69*, 677–681.

Dickinson, A., Hall, G., & Mackintosh, N. J. Surprise and the attenuation of blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, 1976, *2*, 213–222.

Dickinson, A., & Mackintosh, N. J. Reinforcer specificity in the enhancement of conditioning by posttrial surprise. *Journal of Experimental Psychology: Animal Behavior Processes*, 1979, *5*, 162–177.

Dunham, P. J. Punishment: Method and theory. *Psychological Review*, 1971, *78*, 58–70.

Dunham, P. J., Mariner, A., & Adams, H. Enhancement of off-key pecking by on-key punishment. *Journal of the Experimental Analysis of Behavior*, 1969, *1*, 156–166.

Ellias, S. A., & Grossberg, S. Pattern formation, contrast control, and oscillations in the short term memory of shunting on-center off-surround networks. *Biological Cybernetics*, 1975, *20*, 69–98.

Epstein, A. N., Kissileff, H. R., & Stellar, E. (Eds.). *The neuropsychology of thirst: New findings and advances in concepts*. Washington, D.C.: V. H. Winston, 1973.

Estes, W. K. Outline of a theory of punishment. In B. A. Campbell & R. M. Church (Ed.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Estes, W. K., & Skinner, B. F. Some quantitative properties of anxiety. *Journal of Experimental Psychology*, 1941, *29*, 390–400.

Freeman, W. J. *Mass action in the nervous system*. New York: Academic Press, 1975.

Freeman, W. J. EEG analysis gives model of neuronal template-matching mechanism for sensory search with olfactory bulb. *Biological Cybernetics*, 1980, *35*, 221–234.

Freeman, W. J. A neural mechanism for generalization over equivalent stimuli in the olfactory system. In S. Grossberg (Ed.), *Mathematical psychology and psychophysiology*. Providence, R.I.: American Mathematical Society, 1981.

Frey, P. W., & Sears, R. J. Model of conditioning incorporating the Rescorla–Wagner associative axiom, a dynamic attention rule, and a catastrophe rule. *Psychological Review*, 1978, *85*, 321–340.

Friedhoff, A. J. (Ed.). *Catecholamines and behavior: Vol. 1. Basic neurobiology*. New York: Plenum Press, 1975.(a)

Friedhoff, A. J. (Ed.). *Catecholamines and behavior: Vol. 2. Neuropsychopharmacology*. New York: Plenum Press, 1975.(b)

Gabriel, M., Foster, K., Orona, E., Saltwick, S. E., & Stanton, M. Neuronal activity of cingulate cortex, anteroventral thalamus and hippocampal formation in discriminative conditioning: Encoding and extraction of the significance of conditional stimuli. *Progress in Psychobiology and Physiological Psychology*, 1980, *9*, 125–231.

Gibson, J. J. *The ecological approach to visual perception*. Boston: Houghton Mifflin, 1979.

Grossberg, S. *The theory of embedding fields with applications to psychology and neurophysiology*. New York: Rockefeller Institute for Medical Research, 1964.

Grossberg, S. Nonlinear difference-differential equations in prediction and learning theory. *Proceedings of the National Academy of Sciences*, 1967, *58*, 1329–1334.

Grossberg, S. Some physiological and biochemical consequences of psychological postulates. *Proceedings of the National Academy of Sciences*, 1968, *60*, 758–765.

Grossberg, S. Embedding fields: A theory of learning with physiological implications. *Journal of Mathematical Psychology*, 1969, *6*, 209–239.(a)

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

Grossberg, S. Some networks that can learn, remember, and reproduce any number of complicated space-time patterns, I. *Journal of Mathematics and Mechanics*, 1969, *19*, 53–91.(c)

Grossberg, S. On the dynamics of operant conditioning. *Journal of Theoretical Biology*, 1971, *33*, 225–255.

Grossberg, S. A neural theory of punishment and avoidance, I. Qualitative theory. *Mathematical Biosciences*, 1972, *15*, 39–67.(a)

Grossberg, S. A neural theory of punishment and avoidance, II. Quantitative theory. *Mathematical Biosciences*, 1972, *15*, 253–285.(b)

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

Grossberg, S. Contour enhancement, short-term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, 1973, *52*, 217–257.

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

Grossberg, S. A neural model of attention, reinforcement, and discrimination learning. *International Review of Neurobiology*, 1975, *18*, 263–327.

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

Grossberg, S. Adaptive pattern classification and uni-

versal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, 1976, *23*, 187–202.(b)

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

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

Grossberg, S. How does a brain build a cognitive code? *Psychological Review*, 1980, *87*, 1–51.

Grossberg, S. Adaptive resonance in development, perception, and cognition. In S. Grossberg (Ed.), *Mathematical psychology and psychophysiology*. Providence, R.I.: American Mathematical Society, 1981.(a)

Grossberg, S. Psychophysiological substrates of schedule interactions and behavioral contrast. In S. Grossberg (Ed.), *Mathematical psychology and psychophysiology*. Providence, R.I.: American Mathematical Society, 1981.(b)

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

Grossberg, S. *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Boston: Reidel Press, 1982.(b)

Grossberg, S., & Levine, D. S. Some developmental and attentional biases in contrast enhancement and short-term memory of recurrent neural networks. *Journal of Theoretical Biology*, 1975, *53*, 341–380.

Grossberg, S., & Pepe, J. Spiking threshold and overarousal effects in serial learning. *Journal of Statistical Physics*, 1971, *3*, 95–125.

Groves, P. M., Young, S. J., & Wilson, C. J. Nigrostriatal relations and the mechanisms of action of amphetamine. In L. L. Butcher (Ed.), *Cholinergic-monoaminergic interactions in the brain*. New York: Academic Press, 1978.

Hall, G., & Pearce, J. M. Latent inhibition of a CS during CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, 1979, *5*, 31–42.

Hammond, L. J. Retardation of fear acquisition by a previously inhibitory CS. *Journal of Comparative and Physiological Psychology*, 1968, *66*, 756–758.

Hebb, D. O. *The organization of behavior*. New York: Wiley, 1949.

Hebb, D. O. Drives and the C.N.S. (conceptual nervous system). *Psychological Review*, 1955, *62*, 243–254.

Held, R. Exposure-history as a factor in maintaining stability of perception and coordination. *Journal of Nervous and Mental Diseases*, 1961, *132*, 26–32.

Held, R. Dissociation of visual functions by deprivation and rearrangement. *Psychologische Forschung*, 1967, *31*, 338–348.

Held, R., & Hein, A. Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 1963, *56*, 872–876.

Helmholtz, H. von. *Handbuch der physiologischen optik* (1st ed.). Leipzig, German Democratic Republic: Voss, 1866.

Helmholtz, H. von. *Physiological optics* (Vol. 2) (J. P. Southall, Ed.). New York: Dover, 1962. (Originally published, 1866.)

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

Honig, W. K. Attention and the modulation of stimulus control. In D. I. Mostofsky (Ed.), *Attention: Contemporary theory and analysis*. New York: Appleton-Century-Crofts, 1970.

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

Janowitz, H. D., Hanson, M. E., & Grossman, M. I. Effect of intravenously administered glucose on food intake in the dog. *American Journal of Physiology*, 1949, *156*, 87–91.

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

John, E. R. *Mechanisms of memory*. New York: Academic Press, 1967.

Kaczmarek, L. K., & Babloyantz, A. Spatiotemporal patterns in epileptic seizures. *Biological Cybernetics*, 1977, *26*, 199–208.

Kamin, L. J. "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the Prediction of Behavior: Aversive stimulation*. Miami: University of Miami Press, 1968.

Kamin, L. J. Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Keehn, J. D. Effect of shock duration on Sidman avoidance response rates. *Psychological Reports*, 1963, *13*, 852.

Kremer, E. F. Effect of posttrial episodes on conditioning in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, 1979, *5*, 130–141.

Lantz, A. E. Effects of number of trials, interstimulus interval and dishabituation during CS habituation on subsequent conditioning in a CER paradigm. *Animal Learning and Behavior*, 1973, *1*, 273–277.

Le Magnen, J. Regulation of food intake. In F. Reichsman (Ed.), *Hunger and satiety in health and disease*. Basel, Switzerland: S. Karger, 1972.

Levine, D. S., & Grossberg, S. Visual illusions in neural networks: Line neutralization, tilt aftereffect, and angle expansion. *Journal of Theoretical Biology*, 1976, *61*, 477–504.

Lubow, R. E., Rifkin, B., & Alek, M. The context effect: The relationship between stimulus preexposure and environmental preexposure determines subsequent learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 1976, *2*, 38–47.

Macchi, G., & Rinvik, E. Thalamo-telencephalic cir-

cuits: A neuroanatomical survey. In A. Remond (Ed.), *Handbook of electroencephalography and clinical neurophysiology* (Vol. 2, Pt. A). Amsterdam: Elsevier, 1976.

Mackintosh, N. J. An analysis of overshadowing and blocking. *Quarterly Journal of Experimental Psychology*, 1971, *23*, 118–125.

Mackintosh, N. J. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 1975, *82*, 276–298.

Mackintosh, N. J. Overshadowing and stimulus intensity. *Animal Learning and Behavior*, 1976, *4*, 186–192.

Mackintosh, N. J., Bygrave, D. J., & Picton, B. M. B. Locus of the effect of a surprising reinforcer in the attenuation of blocking. *Quarterly Journal of Experimental Psychology*, 1977, *29*, 327–336.

Mackintosh, N. J., & Reese, B. One-trial overshadowing. *Quarterly Journal of Experimental Psychology*, 1979, *31*, 519–526.

Maier, S. F., Seligman, M. E. P., & Solomon, R. L. Pavlovian fear conditioning and learned helplessness effects on escape and avoidance behavior of (a) the CS-US contingency and (b) the independence of the US and voluntary responding. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Masterson, F. A. Is termination of a warning signal an effective reward for the rat? *Journal of Comparative and Physiological Psychology*, 1970, *72*, 471–475.

McAllister, W. R., & McAllister, D. E. Behavioral measurement of conditioned fear. In F. R. Brush (Ed.), *Aversive conditioning and learning*. New York: Academic Press, 1971.

Miller, N. E. Learning resistance to pain and fear: Effects of overlearning, exposure, and rewarded exposure in context. *Journal of Experimental Psychology*, 1960, *60*, 137–145.

Morrell, F. Electrophysiological contributions to the neural basis of learning. *Physiological Review*, 1961, *41*, 443–494.

Näätänen, R., Hukkanen, S., & Järvilechto, T. Magnitude of stimulus deviance and brain potentials. In H. H. Kornhuber & L. Deecke (Eds.), *Progress in brain research: Vol. 54. Motivation, motor and sensory processes of the brain*. New York: Elsevier, 1980.

Newman, F. L., & Baron, M. R. Stimulus generalization along the dimension of angularity: A comparison of training procedures. *Journal of Comparative and Physiological Psychology*, 1965, *60*, 59–63.

O'Keefe, J. O., & Nadel, L. *The hippocampus as a cognitive map*. Oxford, England: Clarendon Press, 1978.

Pearce, J. M., & Hall, G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 1980, *87*, 532–552.

Rescorla, R. A. Establishment of a positive reinforcer through contrast with shock. *Journal of Comparative and Physiological Psychology*, 1969, *67*, 260–263.

Rescorla, R. A., & LoLordo, V. M. Inhibition of avoidance behavior. *Journal of Comparative and Physiological Psychology*, 1965, *59*, 406–412.

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

Routtenberg, A. The two-arousal hypothesis: Reticular formation and limbic system. *Psychological Review*, 1968, *75*, 51–80.

Schwartz, E. L. Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research*, 1980, *20*, 645–669.

Seligman, M. E. P., & Hager, J. L. (Eds.). *Biological boundaries of learning*. New York: Appleton-Century-Crofts, 1972.

Seligman, M. E. P., & Johnston, J. C. A cognitive theory of avoidance learning. In F. J. McGuigan & D. B. Lumsden (Eds.), *Contemporary approaches to conditioning and learning*. Washington, D.C.: V. H. Winston, 1973.

Seward, J. P., & Proctor, D. M. Performance as a function of drive, reward, and habit strength. *American Journal of Psychology*, 1960, *73*, 448–453.

Seward, J. P., Shea, R. A., & Davenport, R. H. Further evidence for the interaction of drive and reward. *American Journal of Psychology*, 1960, *73*, 370–379.

Seward, J. P., Shea, R. A., & Elkind, D. Evidence for the interaction of drive and reward. *American Journal of Psychology*, 1958, *71*, 404–407.

Sharma, K. N., Anand, B. K., Dua, S., & Singh, B. Role of stomach in regulation of activities of hypothalamic feeding centers. *American Journal of Physiology*, 1961, *201*, 593–598.

Singer, W. Control of thalamic transmission by corticofugal and ascending reticular pathways in the visual system. *Physiological Review*, 1977, *57*, 386–420.

Singer, W. Central-core control of visual-cortex functions. In F. O. Schmitt et al. (Eds.), *Neurosciences fourth study program*. Cambridge, Mass.: MIT Press, 1979.

Solomon, R. L., Kamin, L. J., & Wynne, L. C. Traumatic avoidance learning: The outcomes of several extinction procedures with dogs. *Journal of Abnormal and Social Psychology*, 1953, *48*, 291–302.

Strouthes, A. Effects of CS-onset, UCS-termination delay, UCS duration, CS-onset interval, and number of CS-UCS pairings on conditioned fear response. *Journal of Experimental Psychology*, 1965, *69*, 287–291.

Sutton, R. S., & Barto, A. G. Toward a modern theory of adaptive networks: Expectation prediction. *Psychological Review*, 1981, *88*, 135–170.

Tsumoto, T., Creutzfeldt, O. D., & Legendy, C. R. Functional organization of the corticofugal system from visual cortex to lateral geniculate body of the cat. *Experimental Brain Research*, 1976, *25*, 291–306.

Ullman, S. Against direction perception. *Behavioral and Brain Sciences*, 1980, *3*, 373–381.

Wagner, A. R. Frustrative nonreward: A variety of punishment. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Wagner, A. R. Priming in STM: An information pro-

cessing mechanism for self-generated or retrieval-generated depression in performance. In T. J. Tighe & R. N. Leaton (Eds.), *Habituation: Perspectives from child development, animal behavior, and neurophysiology.* Hillsdale, N.J.: Erlbaum, 1976.

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

Wagner, A. R., & Rescorla, R. A. Inhibition in Pavlovian conditioning: Application of a theory. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning.* New York: Academic Press, 1972.

Wagner, A. R., Rudy, J. W., & Whitlow, J. W. Rehearsal in animal conditioning. *Journal of Experimental Psychology,* 1973, *97,* 407–426.

Wallach, H., & Karsh, E. B. The modification of stereoscopic depth-perception and the kinetic depth-effect. *American Journal of Psychology,* 1963, *76,* 429–435.(a)

Wallach, H., & Karsh, E. B. Why the modification of stereoscopic depth-perception is so rapid. *American Journal of Psychology,* 1963, *76,* 413–420.(b)

Wallach, H., Moore, M. E., & Davidson, L. Modification of stereoscopic depth-perception. *American Journal of Psychology,* 1963, *76,* 191–204.

Weisman, R. G., & Litner, J. S. The course of Pavlovian extinction and inhibition of fear in rats. *Journal of Comparative and Physiological Psychology,* 1969, *69,* 667–672.

Received October 29, 1981
Revision received February 7, 1982 ■