NEURAL SUBSTRATES OF ADAPTIVELY TIMED REINFORCEMENT, RECOGNITION, AND MOTOR LEARNING

Stephen Grossberg

May 1995
Revised: October 1995

Technical Report CAS/CNS-95-014

Boston University Center for Adaptive Systems and
Department of Cognitive and Neural Systems
111 Cummington Street
Boston, MA 02215

NEURAL SUBSTRATES OF ADAPTIVELY TIMED REINFORCEMENT, RECOGNITION, AND MOTOR LEARNING

Stephen Grossberg†
Department of Cognitive and Neural Systems
and
Center for Adaptive Systems
Boston University
111 Cummington Street
Boston, MA 02215

May, 1995

Revised: October, 1995

Technical Report CAS/CNS-TR-95-014
Boston, MA: Boston University

To appear in
Models of Action: Mechanisms for Adaptive Behavior
C. Wynne and J.E.R. Staddon (Eds.)
Hillsdale, NJ: Erlbaum Associates, 1995

---

† Supported in part by the Advanced Research Projects Agency (ONR N00014-92-J-4015) and the Office of Naval Research (ONR N00014-92-J-1309, ONR N00014-95-1-0657, and ONR N00014-95-1-0409).

Acknowledgments: We wish to thank Cynthia E. Bradford and Diana Meyers for their valuable assistance in the preparation of the manuscript.

Introduction

The concepts of declarative memory and procedural memory have been used to distinguish two distinct types of learning. A neural network architecture is described that suggests how such memory processes work together as recognition learning, reinforcement learning, and sensory-motor learning all take place together during adaptive behaviors. To coordinate these processes, the hippocampal formation and cerebellum each contain circuits that learn to adaptively time their outputs. Within the architecture, hippocampal timing helps to maintain attention on motivationally salient goal objects during variable task-related delays, and cerebellar timing controls the release of conditioned responses. This property is part of the model's description of how conditionable cognitive-emotional interactions focus attention on motivationally valued cues, and how this process breaks down due to hippocampal ablation. The architecture also suggests how the hippocampal mechanisms that help to rapidly draw attention to salient cues could prematurely release motor commands if the release of these commands were not adaptively timed by the cerebellum.

The model hippocampal system modulates cortical recognition learning without actually encoding the representational information that the cortex encodes. These properties avoid the difficulties faced by several models that propose a direct hippocampal role in recognition learning. Learning within the model hippocampal system controls adaptive timing and spatial orientation. Model properties hereby clarify how hippocampal ablations cause amnesic symptoms and difficulties with tasks which combine task delays, novelty detection, and attention towards goal objects amid distractions. When these model recognition, reinforcement, sensory-motor, and timing processes work together, they suggest how the brain can accomplish conditioning of multiple sensory events to delayed rewards, as during serial compound conditioning.

The chapter is divided into two parts. Part I provides an intuitive introduction to the architecture's main design principles and mechanisms. Part II summarizes mathematical equations and simulations that illustrate how the architecture achieves adaptively timed conditioning and attention.

PART I

How Do Processes of Recognition, Reward, and Action Interact?

A central problem in the behavioral and cognitive neurosciences concerns how humans and other animals learn to recognize objects, to predict and attend to their rewarding or punishing consequences, and to perform appropriately timed actions capable of realizing or avoiding these consequences. Multiple brain regions participate in these processes, including inferotemporal cortex, amygdala, hippocampal formation, and cerebellum. The complexity of these processes has led to the development of neural models that might shed light on their cellular and network properties. A neural architecture is described herein to suggest why both the hippocampus and the cerebellum contain circuits that are specialized for adaptive timing. Although the two timing circuits may share cellular and circuit properties, the architecture predicts that they carry out distinct functional roles during the learning and memory processes that subserve recognition and movement tasks.

These distinct roles are used to clarify several of the conceptual dichotomies that have been useful in research about normal and amnesic learning and memory. One such dichotomy

concerns the distinctions between declarative memory and procedural memory, knowing that and knowing how, memory and habit, or memory with record and memory without record (Bruner, 1969; Mishkin, 1982, 1993; Ryle, 1949; Squire and Cohen, 1984). The amnesic patient HM exemplified this distinction by learning and remembering motor skills like assembly of the Tower of Hanoi without being able to recall having done so (Bruner, 1969; Cohen and Squire, 1980; Mishkin, 1982; Ryle, 1949; Scoville and Milner, 1957; Squire and Cohen, 1984). HM's surgical lesion included extensive parts of the hippocampal formation and amygdala. Subsequent animal studies have shown that damage to the hippocampal formation (Ammon's horn, dentate gyrus, subiculum, fornix) and the parahippocampal region (entorhinal, perirhinal, and parahippocampal cortices) can reproduce analogous amnesic symptoms (Mishkin, 1978; Squire and Zola-Morgan, 1991). These results implicate this aggregate hippocampal system in the processes that regulate declarative memory, or "knowing that". Such processes support a competence for learning recognition categories and being able to flexibly access them in a task-specific way (Eichenbaum, Otto, and Cohen, 1994).

A parallel line of research has implicated the cerebellum in the processing of procedural memory, or "knowing how". The cerebellum is an essential circuit for conditioning discrete adaptive responses during eye movements, arm movements, nictitating membrane movements, and jaw movements (Ebner and Bloedel, 1981; Gilbert and Thach, 1977; Ito, 1984; Lisberger, 1988; Optican and Robinson, 1980; Thompson, 1988; Thompson *et al.*, 1984, 1987). Models of cerebellar learning have been developed over the years to help explain these motor conditioning data (Albus, 1971; Bullock, Fiala, and Grossberg, 1994; Fujita, 1982a, 1982b; Grossberg, 1969a, 1969b; Grossberg and Kuperstein, 1986; Ito, 1984; Lisberger, 1988; Marr, 1969).

A third line of research on learning and memory concerns cognitive-emotional interactions, including how a conditioned stimulus (CS) such as a tone or light, when paired with an unconditioned stimulus (US) such as a shock, can learn to generate conditioned responses (CR), such as fear or limb withdrawal, that were originally elicited only by the US. Such learning is optimal at a range of positive interstimulus intervals (ISI) that are characteristic of the animal and the task, and is greatly attenuated at zero ISI and long ISIs (Smith, 1968). Although the amygdala has been identified as a primary site in the expression of emotion and stimulus-reward association (Aggleton, 1993), the hippocampal formation has also been implicated in the processing of cognitive-emotional interactions. In particular, Thompson *et al.* (1987) distinguished two types of learning that go on during conditioning of the rabbit NMR: "conditioned fear" learning linked to the hippocampus and "learning of the discrete adaptive response" within the cerebellum (p. 82). In addition, removal of the hippocampal formation greatly attenuates attentional blocking (Rickert, Bennett, Lane, and French, 1978; Schmajuk, Spear, and Isaacson, 1983; Solomon, 1977). Blocking is the process whereby conditioning of a cue CS_1 to a US prevents a second cue CS_2 from being conditioned to US when it is later presented before US as part of a simultaneous CS_1 + CS_2 stimulus compound. Much experimental and theoretical work has suggested that CS_2 loses its ability to be conditioned to US because it is an irrelevant cue that predicts no more about the US than does CS_1 when presented alone (Grossberg, 1975, 1982; Kamin, 1969). Blocking enables a learning subject to attend selectively to relevant cues.

The present chapter synthesizes, into a single neural architecture, models that have been developed to explain data from each of these three areas. This synthesis clarifies

how the various models work together to control behavior. In particular, it suggests why both the cerebellum and the hippocampal system may need adaptive timing circuits for their normal functioning. We suggest that the hippocampal mechanisms that help to rapidly draw attention to salient cues could prematurely release motor commands were these commands not adaptively timed by the cerebellum. To reach such conclusions as efficiently as possible, this part of the article provides just enough information about the component models to understand how they can work together to explain key data. Mathematical equations of some key model processes are described in Part II. Others are developed in detail in other articles that are cited below.

The chapter is devoted to an exposition of just one theory for several reasons. One reason is space limitations. Another is that no other theory of which we are aware has explained such a large data base or articulated the design principles that support this explanatory range. Some comparisons with other models are found in Grossberg and Merrill (1995).

Multiple Roles for the Hippocampal System?

Why should a single, albeit complex, brain region like the hippocampal system be involved in so many processes: recognition learning, reinforcement learning, and motivated attention? A clue is provided by neural data and models about how each of these processes works. In particular, both recognition learning and reinforcement learning are regulated by a matching process whereby bottom-up stimuli from the outside world are matched against top-down learned expectations to determine whether attentive learning or memory search will occur. The unblocking paradigm illustrates this matching process for the case of reinforcement learning (Kamin, 1969). The unblocking paradigm is a variant of the blocking paradigm in which the US changes intensity in the two learning episodes. Thus if CS_1 is followed by one US intensity (US_1), and the compound stimulus $CS_1 + CS_2$ is followed by a different US intensity (US_2), then CS_2 can become conditioned to the US, unlike in the blocking paradigm, and does so with an emotional valence that depends upon the sign of the difference $US_1 - US_2$ between US_1 and US_2 (Kamin, 1969). The mismatch between the actual intensity US_2 and the expected intensity US_1 triggers a memory search that attentionally "unblocks" the representation of CS_2 that is stored in short term memory, and enables it to learn to predict the change in US intensity (Grossberg, 1975). This memory search helps to focus attention upon that subset of sensory cues that predicts motivationally salient outcomes in a given context, and to block those that do not.

Recognition learning is accomplished by interactions between inferotemporal cortex (IT) and hippocampal formation, among other brain areas (Desimone, 1991; Desimone and Ungerleider, 1989; Eichenbaum, Otto, and Cohen, 1994; Gochin, Miller, Gross, and Gerstein, 1991; Harries and Perrett, 1991; Mishkin, 1978, 1982; Mishkin and Appenzeller, 1987; Perrett, Mistlin, and Chitty, 1987; Schwartz, Desimone, Albright, and Gross, 1983; Squire and Zola-Morgan, 1991). These interactions include the matching process that modulates the course of recognition learning in IT cortex and the course of reinforcement learning in thalamocortico-amygdala circuits. Some models are analysed below of how these recognition and reinforcement learning circuits interact with motor learning circuits. It is shown that the behavioral success of this interaction requires both types of circuits to be adaptively timed.

Self-Organizing Feature Maps and Adaptive Resonance

The first type of model results from an analysis of how humans and animals rapidly learn to categorize and name events and their contexts in real time. These Adaptive Resonance Theory (ART) models have been used to help explain and predict a large body of cognitive and neural data about recognition learning, recall, attention, priming, and memory search (Carpenter and Grossberg, 1992, 1993; Grossberg, 1982, 1987, 1988a). ART systems realize this synthesis by incorporating mechanisms that solve a fundamental problem about learning and memory that is called the *stability-plasticity dilemma.* An adequate self-organizing recognition system must be capable of plasticity in order to rapidly learn about significant new events, yet its memory must also remain stable in response to irrelevant or often repeated events. Thus we can learn to recognize many new faces without risking the unselective forgetting of our parents' faces. In ART, interactions between an attentional subsystem and an orienting subsystem, or novelty detector, self-stabilize the learning process as the network becomes familiar with an environment by categorizing the information within it in a way that leads to behavioral success (Grossberg, 1980).

Learning takes place in the attentional subsystem. Its processes include activation of short term memory (STM) traces, incorporation through learning of STM information into a longer-lasting long term memory (LTM) traces, and interactions between pathways that carry specific information with nonspecific pathways that modulate the specific pathways. These interactions between specific STM and LTM processes and nonspecific modulatory processes regulate the stability-plasticity balance during normal learning, as follows.

Figure 1

The attentional subsystem undergoes both bottom-up learning and top-down learning between processing levels such as those denoted by $\mathcal{F}_1$ and $\mathcal{F}_2$ in Figure 1. Level $\mathcal{F}_1$ contains a network of nodes, or cell populations, each of which is activated by a particular combination of sensory features. Level $\mathcal{F}_2$ contains a network of nodes that represent recognition codes, or categories, which are selectively activated by the activation patterns across $\mathcal{F}_1$. Each $\mathcal{F}_1$ node sends output signals to a subset of $\mathcal{F}_2$ nodes. Each $\mathcal{F}_2$ node thus receives inputs from many $\mathcal{F}_1$ nodes. The thick pathway from $\mathcal{F}_1$ to $\mathcal{F}_2$ in Figure 1 represents the array of diverging and converging pathways, for simplicity. Learning takes place at the synapses denoted by semicircular endings in the $\mathcal{F}_1 \rightarrow \mathcal{F}_2$ pathways. Pathways that end in arrowheads do not undergo learning. This bottom-up learning enables $\mathcal{F}_2$ nodes to become selectively tuned to particular combinations of activation patterns across $\mathcal{F}_1$ by changing their LTM traces.

Why is bottom-up learning insufficient in a system that can autonomously solve the stability-plasticity dilemma? This analysis was carried out in that part of the ART model that combines bottom-up associative learning and lateral inhibition for purposes of learned categorization. This type of model is often called a self-organizing feature map, competitive learning, or learned vector quantization. In such a model, as shown in Figure 2A, an input pattern registers itself as a pattern of activity, or STM, across the feature detectors of level $\mathcal{F}_1$. Each $\mathcal{F}_1$ output signal is multiplied or gated, by the adaptive weight, or LTM trace, in its respective pathway. All these LTM-gated inputs are added up at their target $\mathcal{F}_2$ nodes. Competitive interactions, mediated by lateral inhibition within $\mathcal{F}_2$, contrast-enhance this input pattern. Even if many $\mathcal{F}_2$ nodes may receive inputs from $\mathcal{F}_1$, lateral inhibition acts to cause a much smaller set of $\mathcal{F}_2$ nodes to store their activation in STM.

It is useful to think of all the STM signals that converge on an $\mathcal{F}_2$ node as an STM pattern, or vector. Likewise, all the LTM traces that multiply these signals on their way to a prescribed $\mathcal{F}_2$ node form an LTM vector. The operation of adding up the LTM-gated signals at each $\mathcal{F}_2$ node is called the inner product, or dot product, of the two vectors. It measures how similar the two vectors are, and increases as a function of their similarity. The LTM traces thereby *filter* the STM signal pattern and generate larger inputs to those $\mathcal{F}_2$ nodes whose LTM patterns are most similar to the STM pattern.

As noted above, the lateral inhibition among $\mathcal{F}_2$ nodes selects just a few of the more active $\mathcal{F}_2$ nodes for STM storage. This contrast-enhancing operation enables many input patterns at $\mathcal{F}_1$ that share similar input features to be classified by a small set of $\mathcal{F}_2$ nodes. The $\mathcal{F}_2$ nodes hereby become category nodes that are capable of classifying the inputs to $\mathcal{F}_1$.

Figure 2

In a self-organizing feature map, only the $\mathcal{F}_2$ nodes that win the contrast-enhancing competition and store their activity in STM can influence the learning process. STM activity at the winning $\mathcal{F}_2$ nodes selectively opens a learning gate at the LTM traces that abut these nodes. These LTM traces can then approach, or track, the input signals in their pathways, a process called steepest descent. This learning law is thus often called gated steepest descent, or instar learning. In its simplest form, this learning law can be expressed by the equation

$$\frac{d}{dt}w_{ij} = f(x_j)(-w_{ij} + S_i), \tag{1}$$

where $\frac{d}{dt}w_{ij}$ is the time rate of change of the LTM trace, or adaptive weight, w_{ij} from the i^{th} $\mathcal{F}_1$ node to the j^{th} $\mathcal{F}_2$ node, $f(x_j)$ is the learning gating signal that becomes positive only if the postsynaptic activity, or potential, x_j of the j^{th} $\mathcal{F}_2$ node becomes sufficiently large, and S_i is the i^{th} bottom-up signal. This learning rule was introduced into neural network models in Grossberg (1969a) and is the learning rule that was used to introduce ART (Grossberg, 1976b). While tracking the signals in its pathway, such an LTM trace w_{ij} can either increase (if the signal S_i is large) or decrease (if the signal S_i is small). It thus combines Hebbian and anti-Hebbian learning properties in a way that has been used to model neurophysiological data about hippocampal LTP and LTD (Artola and Singer, 1993; Levy, 1985; Levy and Desmond, 1985) and adaptive tuning of cortical feature detectors during the visual critical period (Rauschecker and Singer, 1979; Singer, 1983).

In particular, as Table 1 shows, significant postsynaptic activity, mediated by the gating signal $f(x_j)$, is needed to cause any change in w_{ij}. If this modulatory gate opens, then w_{ij} may increase or decrease, depending upon the relative size of S_i. Since S_i, in turn, may influence the amount of postsynaptic activity x_j via the presynaptic signal $S_i w_{ij}$, various secondary effects can occur that are beyond the scope of this discussion (but see Carpenter and Grossberg, 1990). It is perhaps worth noting, however, that an early prediction (Grossberg, 1968b, 1969b, 1974) suggested that synaptic learning would be mediated by a postsynaptic process of protein synthesis and receptor sensitization that controls a coordinated presynaptic process of transmitter production. The postsynaptic signal process was predicted to be triggered by an inward Ca^{++} current that is antagonistic to Mg^{++}. Coordinated presynaptic and postsynaptic changes were predicted to depend upon the inward Ca^{++} current in synergy with an inward Na^{+} current and an outward K^{+} current. Similar

concepts have been used, in greatly elaborated form, to explain recent data about LTP and LTD; e.g., see Artola and Singer (1993) and Kuno (1995). Gated steepest descent learning may thus be viewed as a first approximation to a much more complex cascade of biochemical events.

Table 1

The net effect of such learning is to train the LTM vectors of the winning $\mathcal{F}_2$ category nodes to become more similar to the STM patterns that they filter. As a result, the winning $\mathcal{F}_2$ categories sharpen their tuning curves to respond more selectively to the STM patterns that they have experienced.

Self-organizing feature map models were introduced and computationally characterized in Malsburg (1973) and Grossberg (1976a, 1978). These models were subsequently applied and further developed by many authors, notably Kohonen (1984). They exhibit many useful properties, especially if not too many input patterns, or clusters of input patterns, perturb level $\mathcal{F}_1$ relative to the number of categorizing nodes in level $\mathcal{F}_2$. Grossberg (1976a) proved under these sparse environmental conditions that category learning is stable, with LTM traces that track the statistics of the environment, are self-normalizing, and oscillate a minimum number of times. Also, the $\mathcal{F}_2$ category selection rule, like a Bayesian classifier, tends to minimize error.

It was also proved, however, that under more general environmental conditions, learning becomes unstable and subject to catastrophic forgetting. Such a model could forget the faces of your parents while learning a new face. This memory instability is due to basic properties of associative learning and lateral inhibition. Although a gradual switching off of plasticity can partially overcome the problem, such a mechanism cannot work in a learning system whose plasticity is maintained throughout adulthood. These results put into sharp focus the problem of how the brain dynamically self-stabilizes its memory while remaining open to new experiences throughout life, a topic that has attracted increasing interest (Kandel and O'Dell, 1992). An analysis of this instability, together with data about categorization, conditioning, and attention, led to the introduction of ART models that self-stabilize the memory of self-organizing feature maps in response to an arbitrary stream of input patterns (Grossberg, 1976b).

The Link Between Top-Down Matching, Hypothesis Testing, and Stable Learning

In an ART model, learning does not occur when some winning $\mathcal{F}_2$ activities are stored in STM. Instead activation of $\mathcal{F}_2$ nodes may be interpreted as "making a hypothesis" about an input at $\mathcal{F}_1$. When $\mathcal{F}_2$ is activated, it quickly generates an output pattern that is transmitted along the top-down adaptive pathways from $\mathcal{F}_2$ to $\mathcal{F}_1$. These top-down signals are multiplied in their respective pathways by LTM traces at the semicircular synaptic knobs of Figure 2B. The LTM-gated signals from all the active $\mathcal{F}_2$ nodes are added to generate the total top-down feedback pattern from $\mathcal{F}_2$ to $\mathcal{F}_1$. This pattern plays the role of a learned expectation. Activation of this expectation may be interpreted as "testing the hypothesis", or "reading out the prototype", of the active $\mathcal{F}_2$ category. As shown in Figure 2B, ART networks are designed to match the "expected prototype" of the category against the bottom-up input pattern, or exemplar, to $\mathcal{F}_1$. Nodes that are activated by this exemplar are suppressed if they

do not correspond to large LTM traces in the top-down prototype pattern. The resultant $\mathcal{F}_1$ pattern encodes the cluster of input features that the network deems relevant to the hypothesis based upon its past experience. This resultant activity pattern, called $\mathbf{X}^*$ in Figure 2B, encodes the pattern of features to which the network "pays attention".

If the expectation is close enough to the input exemplar, then a state of *resonance* develops as the attentional focus takes hold. The pattern $\mathbf{X}^*$ of attended features reactivates the $\mathcal{F}_2$ category $\mathbf{Y}$ which, in turn, reactivates $\mathbf{X}^*$. The network locks into a resonant state through a positive feedback loop that dynamically links, or binds, $\mathbf{X}^*$ with $\mathbf{Y}$. Damasio (1989) has used the term "convergence zones" to describe such a resonant process. Such resonances are capable of binding spatially distributed features into synchronous and coherent states, both in cortico-cortical and thalamocortical feedback networks (Grossberg, 1976b; Grossberg and Somers, 1991).

Neurophysiological data that are consistent with the prediction that ART-like resonances exist between LGN and V1 have recently been reported (Sillito, Jones, Gerstein, and West, 1994). In particular, it was suggested in Grossberg (1980) that top-down corticogeniculate feedback would selectively amplify monocular LGN activations that are consistent with the oriented binocular cortical cells that activate the feedback, while inhibiting LGN cells that are not. In addition, top-down feedback by itself, as in all ART systems, was suggested not to be fully able to activate LGN cells. In support of this prediction, Sillito *et al.* (1994) reported that "cortically induced correlation of relay cell activity produces coherent firing in those groups of relay cells with receptive field alignments appropriate to signal the particular orientation of the moving contour to the cortex ... this increases the gain of the input for feature-linked events detected by the cortex ... the cortico-thalamic input is only strong enough to exert an effect on those LGN cells that are additionally polarized by their retinal input ... the feedback circuit searches for correlations that support the 'hypothesis' represented by a particular pattern of cortical activity" (pp. 479–482). Gove, Grossberg, and Mingolla (1995) have shown how this type of corticogeniculate feedback and resonance can be used as part of a larger model of cortical visual processing to simulate data about brightness perception and illusory contours.

Similar ART matching and resonance rules have been used to explain and predict behavioral and brain data from other task domains. For example, Carpenter and Grossberg (1993) have used ART matching and resonance rules to explain data about visual object recognition and medial temporal amnesia (see below). Govindarajan, Grossberg, Wyse, and Cohen (1994) have used ART matching and resonance rules to simulate auditory psychophysical data about acoustic source segregation when multiple sources harmonically overlap, as during a cocktail party. Grossberg, Boardman, and Cohen (1994) have used ART matching and resonance rules to simulate psychophysical data about variable-rate speech categorization. Grossberg and Stone (1986a) have used such rules to explain data about lexical priming and decision making. Roberts, Aguilar, Bullock, and Grossberg (1994) have used ART matching and resonance rules to explain neural data about multimodal control of saccadic eye movements. Why should similar matching and resonance rules be used in so many brain systems?

ART shows how these matching and resonance rules can be used to help solve the noise-saturation dilemma in any brain system that dynamically adjusts and maintains its parameters to cope with changing environmental conditions throughout life. The matched

resonant state, rather than bottom-up activation, is predicted to drive the learning process. The resonant state persists long enough, at a high enough activity level, to activate the slower learning process; hence the term *adaptive resonance* theory. ART systems learn prototypes, rather than exemplars, because the attended feature vector $\mathbf{X}^*$, rather than the input exemplar itself, is learned. Both the bottom-up LTM traces that tune the category nodes and the top-down LTM traces that filter the learned expectation learn to correlate activation of $\mathcal{F}_2$ nodes with the set of all *attended* $\mathbf{X}^*$ vectors that they have ever experienced. These attended STM vectors assign less STM activity to features in the input vector $\mathbf{I}$ that mismatch the learned top-down prototype $\mathbf{V}$ than to features that match $\mathbf{V}$. Bottom-up activations that are not supported by large top-down LTM traces are hereby suppressed, and hence cannot destabilize the learning process.

Prototype Learning and Exemplar Learning Using Vigilance Control

A similar type of matching by similarity across arrays of features has been used to quantitatively fit categorization data from human subjects (Estes, 1994). Models of this type need to assume that every input exemplar that a subject has ever processed is stored. Such models face formidable problems of memory storage and retrieval, and have not yet been shown capable of real-time autonomous categorization of complex databases. ART models computationally elaborate the idea that humans learn prototypes (Posner and Keele, 1968, 1970), which save greatly on memory resources by allowing many exemplars to be represented by a single category prototype. ART models have also been used for real-time autonomous categorization of complex databases (e.g., Asfour, Carpenter, and Grossberg, 1995; Asfour *et al.*, 1993; Bachelder, Waxman, and Seibert, 1993; Baloch and Waxman, 1991; Bradski and Grossberg, 1994; Carpenter *et al.*, 1992; Carpenter, Grossberg, and Reynolds, 1991, 1995; Carpenter and Ross, 1993; Carpenter and Tan, 1993; Caudell, Smith, Escobedo, and Anderson, 1994; Dubrawski and Crowley, 1994; Gjerdingen, 1990; Goodman *et al.*, 1992; Ham and Han, 1993; Harvey, 1993; Kasperkiewicz, Racz, and Dubrawski, 1994; Keyvan, Durg, and Rabelo, 1993; Metha, Vij, and Rabelo, 1993; Moya, Koch, and Hostetler, 1993; Seibert and Waxman, 1992; Suzuki, Abe, and Ono, 1994; Suzuki, 1995; Wienke, Xie, and Hopke, 1994).

Given that ART systems learn prototypes, how can they also learn to recognize unique experiences, such as a particular view of a friend's face? The prototypes that are learned by ART systems accomplish this by realizing a qualitatively different concept of prototype than that offered by previous models. In particular, ART prototypes form in a way that is designed to conjointly maximize category generalization while minimizing predictive error (Carpenter, Grossberg, and Reynolds, 1991; Carpenter *et al.*, 1992). As a result, ART prototypes can automatically learn individual exemplars when environmental conditions require highly selective discriminations to be made. How the matching process achieves this is discussed below.

First, let us consider what happens if the mismatch between bottom-up and top-down information is too great for a resonance to develop. Then the $\mathcal{F}_2$ category is quickly reset and a memory search for a better category is initiated. This combination of top-down matching, attention focusing, and memory search is what stabilizes ART learning and memory in an arbitrary input environment. The attentional focusing by top-down matching prevents inputs that represent irrelevant features at $\mathcal{F}_1$ from eroding the memory of previously learned

LTM prototypes. In addition, the memory search resets $\mathcal{F}_2$ categories so quickly when their prototype $\mathbf{V}$ mismatches the input vector $\mathbf{I}$ that the more slowly varying LTM traces do not have an opportunity to correlate the attended $\mathcal{F}_1$ activity vector $\mathbf{X}^*$ with them. Conversely, the resonant event, when it does occur, maintains and amplifies the matched STM activities for long enough and at high enough amplitudes for learning to occur in the LTM traces.

Whether or not a resonance occurs depends upon the level of mismatch, or novelty, that the network is prepared to tolerate. Novelty is measured by how well a given exemplar matches the prototype that its presentation evokes. The criterion of an acceptable match is defined by an internally controlled parameter ρ called vigilance (Carpenter and Grossberg, 1987a, 1992). The vigilance parameter is computed in the orienting subsystem $\mathcal{A}$; see Figure 1. Vigilance weighs how similar an input exemplar $\mathbf{I}$ must be to a top-down prototype $\mathbf{V}$ in order for resonance to occur. Resonance occurs if $\rho|\mathbf{I}| - |\mathbf{X}^*| \leq 0$. This inequality says that the $\mathcal{F}_1$ attentional focus $\mathbf{X}^*$ inhibits $\mathcal{A}$ more than the input $\mathbf{I}$ excites it. If $\mathcal{A}$ remains quiet, then an $\mathcal{F}_1 \leftrightarrow \mathcal{F}_2$ resonance can develop and category learning can ensue.

Either a larger value of ρ or a smaller match ratio $|\mathbf{X}^*||\mathbf{I}|^{-1}$ makes it harder to satisfy the resonance inequality. When ρ grows so large or $|\mathbf{X}^*||\mathbf{I}|^{-1}$ is so small that $\rho|\mathbf{I}| - |\mathbf{X}^*| > 0$, then $\mathcal{A}$ generates an arousal burst, or novelty wave, that resets the STM pattern across $\mathcal{F}_2$ before new learning can occur, and initiates a bout of hypothesis testing, or memory search. During search, the orienting subsystem interacts with the attentional subsystem (Figures 2C and 2D) to rapidly reset mismatched categories and to select better $\mathcal{F}_2$ representations with which to categorize novel events at $\mathcal{F}_1$, without risking unselective forgetting of previous knowledge. Search may select a familiar category if its prototype is similar enough to the input to satisfy the resonance criterion. The prototype may then be refined by attentional focusing before learning occurs. If the input is too different from any previously learned prototype, then an uncommitted population of $\mathcal{F}_2$ cells is selected and learning of a new category is initiated.

Because vigilance can vary across learning trials, recognition categories capable of encoding widely differing degrees of generalization or abstraction can be learned by a single ART system. Low vigilance leads to broad generalization and abstract prototypes. High vigilance leads to narrow generalization and to prototypes that represent fewer input exemplars, even a single exemplar. Thus a single ART system may be used, say, to learn abstract prototypes with which to recognize abstract categories of faces and dogs, as well as "exemplar prototypes" with which to recognize individual faces and dogs. A single system can learn both, as the need arises, by increasing vigilance just enough to activate $\mathcal{A}$ if a previous categorization leads to a predictive error (Carpenter *et al.*, 1992; Carpenter, Grossberg, and Reynolds, 1991).

Corticohippocampal Interactions and Medial Temporal Amnesia

As sequences of inputs are practiced over learning trials, the search process eventually converges upon stable categories. It has been mathematically proved (Carpenter and Grossberg, 1987a, 1992) that familiar inputs directly access the category whose prototype provides the globally best match, without requiring a search. This property helps to explain how we can recognize familiar objects so quickly, even though we may know about many other things. Unfamiliar inputs continue to engage the orienting subsystem to trigger memory searches for better categories until they become familiar. New categories can continue to form until

the memory capacity is fully utilized. Memory capacity is determined by the number of category nodes in $\mathcal{F}_2$, which can be chosen to be as large as one desires without degrading system performance.

The process whereby search is automatically disengaged is a form of memory consolidation that emerges from network interactions. This type of "emergent consolidation" does not preclude structural consolidation at individual cells, since the amplified and prolonged activities that subserve a resonance may be a trigger for learning-dependent cellular processes, such as protein synthesis and transmitter production.

The attentional subsystem of ART has been used to model aspects of inferotemporal (IT) cortex, and the orienting subsystem models part of the hippocampal system. The interpretation of ART dynamics in terms of IT cortex led Miller, Li, and Desimone (1991) to successfully test the prediction that cells in monkey IT cortex are reset after each trial in a working memory task. To illustrate the implications of an ART interpretation of IT-hippocampal interactions, Carpenter and Grossberg (1993) have described how a lesion of the ART model's orienting subsystem creates a formal memory disorder with symptoms much like the medial temporal amnesia that is caused in animals and patient HM after hippocampal system lesions. In particular, such a lesion *in vivo* causes unlimited anterograde amnesia; limited retrograde amnesia; failure of consolidation; tendency to learn the first event in a series; abnormal reactions to novelty, including perseverative reactions; normal priming; and normal information processing of familiar events (Cohen, 1984; Graf, Squire, and Mandler, 1984; Lynch, McGaugh, and Weinberger, 1984; Squire and Butters, 1984; Squire and Cohen, 1984; Warrington and Weiskrantz, 1974; Zola-Morgan and Squire, 1990).

Unlimited anterograde amnesia occurs because the network cannot carry out the memory search to learn a new recognition code. Limited retrograde amnesia occurs because familiar events can directly access correct recognition codes. Before events become familiar, memory consolidation occurs which utilizes the orienting subsystem (Figure 1C). This failure of consolidation does not necessarily prevent learning *per se*. Instead, learning influences the first recognition category activated by bottom-up processing, much as "amnesics are particularly strongly wedded to the first response they learn" (Gray, 1982, p. 253). Perseverative reactions can occur because the orienting subsystem cannot reset sensory representations or top-down expectations that may be persistently mismatched by bottom-up cues. The inability to search memory prevents ART from discovering more appropriate stimulus combinations to attend. Normal priming occurs because it is mediated by the attentional subsystem.

Similar behavioral problems have been identified in hippocampectomized monkeys. Gaffan (1985) noted that fornix transection "impairs ability to change an established habit ... in a different set of circumstances that is similar to the first and therefore liable to be confused with it" (p. 94). In ART, a defective orienting subsystem prevents the memory search whereby different representations could be learned for similar events. Pribram (1986) called such a process a "competence for recombinant context-sensitive processing" (p. 362). These ART mechanisms illustrate how memory consolidation and novelty detection may be mediated by the same neural structures (Zola-Morgan and Squire, 1990), why hippocampectomized rats have difficulty orienting to novel cues (O'Keefe and Nadel, 1978), and why there is a progressive reduction in novelty-related hippocampal potentials as learning proceeds in normal rats (Deadwyler, West, and Lynch, 1979; Deadwyler, West, and Robinson, 1981). In ART, the orienting system is automatically disengaged as events become familiar during the

memory consolidation process.

ART properties hereby provide an alternative to the popular hypothesis that the hippocampal formation somehow temporarily stores recognition codes from all sensory modalities before the temporal cortex can more permanently do so (Eichenbaum, Otto, and Cohen, 1994; Marr, 1971; McClelland, McNaughton, and O'Reilly, 1994; Milner, 1989). This hypothesis faces formidable obstacles as soon as one seriously tries to model how such a process could work. For example, how could the hippocampal system rapidly store all the information that one can recall after seeing an exciting movie? McClelland, McNaughton, and O'Reilly (1994) admit that their model cannot do this. In fact, not only is fast learning impossible, but also "the sequential acquisition of new data ... can lead to *catastrophic interferences* with what has previously been learned". Only if learning is slow and carefully interleaved on sufficiently small and regular databases can it occur at all in this type of model. Such a model fails to solve the stability-plasticity dilemma. Grossberg and Merrill (1995) provide a comparative analysis of the ART corticohippocampal model of medial temporal amnesia with alternative amnesia models, both in terms of their explanatory power and their plausibility as neural mechanisms.

A Prediction About How Corticohippocampal Interactions Control the Specificity of Learned Prototypes

The ART conception of temporal-hippocampal interactions suggests the following prediction. Level $\mathcal{F}_2$ properties may be compared with properties of cell activations in inferotemporal cortex (IT) during recognition learning in monkeys. The ability of $\mathcal{F}_2$ nodes to learn categories with different levels of generalization clarifies how some IT cells can exhibit high specificity, such as selectivity to views of particular faces, while other cells respond to broader features of the animal's environment (Desimone and Ungerleider, 1989; Gochin *et al.*, 1991; Harries and Perrett, 1991; Mishkin, 1982; Mishkin and Appenzeller, 1987; Perrett, Mistlin, and Chitty, 1987; Schwartz *et al.*, 1983; Seibert and Waxman, 1991). Moreover, when monkeys are exposed to easy and difficult discriminations (Spitzer, Desimone, and Moran, 1988), "in the difficult condition the animals adopted a stricter internal criterion for discriminating matching from nonmatching stimuli... the animals' internal representations of the stimuli were better separated, independent of the criterion used to discriminate them... increased effort appears to cause enhancement of the responses and sharpened selectivity for attended stimuli" (pp. 339–340). These are also properties of model cells in $\mathcal{F}_2$ due to the role of vigilance control. ART prototypes represent smaller sets of exemplars at higher vigilance levels, so a stricter matching criterion is learned. These exemplars match their finer prototypes better than do exemplars which match a coarser prototype. This better match more strongly activates the corresponding $\mathcal{F}_2$ nodes.

This property suggests that operations which make the novelty-related potentials of the hippocampus more sensitive to input changes may trigger the formation of more selective inferotemporal recognition categories. Can such a correlation between IT discrimination and hippocampal potentials be recorded, say, when monkeys learn easy and difficult discriminations? Conversely, operations that progressively block the expression of hippocampal novelty potentials are suggested to cause learning of coarser recognition categories, with amnesic symptoms as a limiting case.

The conclusion that no learning occurs in the ART orienting system does not force

the theory to deny that some types of learning do occur in the hippocampal system. The model suggests that these learning processes are involved in adaptively timed modulation of reinforcement learning and aspects of spatial orientation, as discussed below.

A Framework for Temporal Learning

Before providing this discussion, it is appropriate to comment upon how an ART-based system could rapidly learn the information in a movie. Such an analysis requires that the processes whereby individual events are recognized and recalled are supplemented by processes involved in the learning and recognition of temporally ordered sequences, or lists, of events. There are many levels on which this class of problems could be approached, and it seems fair to say that no available theory proposes a complete explanation of this competence. On the other hand, the critique of alternative models has been made on the level of their inability to rapidly and stably learn large amounts of information, notably temporally ordered information. This is not a problem in an ART-based system.

A framework for accomplishing this was described in Grossberg (1978) using a combination of ART category learning, working memories, temporal associative learning networks, and predictive feedback within the system. A great deal of work has since been done to further carry out this program. For example, ART-based architectures, called VIEWNET systems, are capable of rapidly and stably learning to recognize 3-D objects by categorizing their 2-D views and learning to associate their 2-D view categories with 3-D object categories that are invariant under changes of familiar 2-D view (Bradski and Grossberg, 1994, 1995). Properties of these 2-D view and 3-D object category nodes may be compared with neural responses from distinct cell populations in monkey inferotemporal cortex (Logothetis *et al.*, 1994).

The 3-D object categories may, in turn, be stored in a working memory (Baddeley, 1986; Grossberg, 1978) that can encode both object representations and their temporal order in STM. This type of working memory is designed so that its contents may rapidly and stably be learned and categorized by another ART network, whose active nodes are said to code list categories. This list categorization process has been proved to retain its stability even as new information continues to be stored in the working memory through time (Bradski, Carpenter, and Grossberg, 1992, 1994; Cohen and Grossberg, 1986, 1987; Grossberg, 1978; Grossberg and Stone, 1986a). Interactions between such a working memory and its list categories have been used to explain data from experiments about the sequential performance of stored motor commands (Boardman and Bullock, 1991; Grossberg and Kuperstein, 1989), about errors in serial item and order recall due to rapid visual attention shifts (Grossberg and Stone, 1986a), about errors and reaction times during lexical priming and episodic memory experiments (Grossberg and Stone, 1986b), and about data concerning word superiority, phonemic restoration, and backward effects on speech perception (Cohen and Grossberg, 1986; Grossberg, 1986). Such a working memory design thus seems to be used in several modalities. This is plausible when one realizes that the design embodies a few simple principles that enable its temporally evolving STM patterns to be stably categorized in LTM.

Temporal cortex provides a likely neural substrate for such a working memory (Goldman-Rakic, 1994). Here, information from multiple sensory modalities converges and may interact with subcortical reward mechanisms to sustain an attentional focus upon salient goals (Gaffan, 1994; Knight, 1994). Can ART systems learn multimodal list categories and focus

attention on predictively successful ones?

Multimodal information distributed across a working memory may indeed be integrated into ART categories (Asfour, 1994; Asfour *et al.*, 1993). Such an ART system, called Fusion ARTMAP, is designed to solve the credit assignment problem of selectively resetting those input channels that are causing predictive errors. In addition, ART models of cognitive-emotional interactions have been described to suggest how attention may be selectively allocated to event categories that have high salience due to prior reinforcement and how less salient events may be attentionally blocked (Grossberg, 1975, 1982, 1984; Grossberg and Levine, 1987; Grossberg and Merrill, 1992); also see below. They have also been used to explain and predict cognitive data about human decision making under risk as a manifestation of cognitive-emotional neural mechanisms (Grossberg and Gutowski, 1987), and to shed some light upon how these cognitive-emotional interactions may break down during mental depression (Grossberg, 1972b, 1984).

The motivationally modulated list categories may, in turn, be recurrently linked together by an associative learning network that helps to predict the categories most likely to occur in a given temporal context. Such networks have been used to model the position-dependent error gradients and learning rates that are observed during human verbal learning and to predict how this process breaks down in schizophrenic subjects (Grossberg, 1969c, 1982b; Grossberg and Pepe, 1970, 1971). Finally, the attended list categories may be used to predict the next images that are expected by the system, a one-to-many process called outstar learning (Grossberg, 1968a, 1978, 1980). One possible anatomical substrate of this type of predictive learning is frontotemporal projections (Gaffan, 1994).

Taken together, these architectural elements may be called a *resonant avalanche*. This name acknowledges the role of resonance in stabilizing the learning process, and of the avalanche of temporal associations in predicting the events that the system next expects to experience. (For a summary of avalanches at different levels of complexity, see Grossberg, 1978.) Although the theory of resonant avalanches has not yet been completely developed, there are enough mathematical, computational, and data simulation results available to conclude that ART systems escape the critique of various other memory models that was proposed above.

Adaptively Timed Cognitive-Emotional and Sensory-Motor Interactions

Let us now return to the question of what sorts of learning are predicted to occur in the hippocampal system by an ART-based model. As in our remarks about fronto-temporal interactions, this discussion will include an analysis of issues concerning reinforcement and temporal processing. The model fronto-temporal interactions that were reviewed above concern a type of *macro*-timing that integrates information across a series of events. The model fronto-temporal-hippocampal interactions now to be discussed consider a type of *micro*-timing that calibrates how long motivated attention may be allocated to a single predicted event.

Some authors (e.g., Eichenbaum, Otto, and Cohen, 1994) have dichotomized the representational properties of hippocampal memory processing–namely, those relating to recognition learning and memory—as being "orthogonal functional properties" from hippocampal temporal processing properties. It is unclear why a single brain structure should combine properties if they are indeed "orthogonal". The adaptive timing model described below sug-

gests how these representational and temporal processes may be linked. The timing model is part of a larger model system that controls how cognitive-emotional and sensory-motor interactions are coordinated, including how classical and instrumental conditioning are adaptively timed and modulated by cognitive recognition processes (Baloch and Waxman, 1991; Grossberg, 1971, 1972b, 1975, 1982a, 1987; Grossberg and Levine, 1987; Grossberg and Merrill, 1992; Grossberg and Schmajuk, 1987).

Figure 3

This cognitive-emotional model suggests that (at least) three types of internal representation interact during conditioning: sensory representations S, drive representations D, and motor representations M (Figure 3). The S representations are categorical thalamocortical representations of external events, including the object recognition categories that are learned by IT cortex and linked to frontal cortex via fronto-temporal interactions. The D representations include hypothalamic and amygdala circuits, at which homeostatic and reinforcing cues converge to generate emotional reactions and motivational decisions. The M representations include cerebellar circuits that control discrete adaptive responses. Three types of learning take place among these representations: S $\rightarrow$ D conditioned reinforcer learning that converts a CS into a reinforcer by pairing activation of its sensory representation S with activation of the drive representation D that receives input from a salient US or other conditioned reinforcer CS; D $\rightarrow$ S incentive motivational learning whereby an activated drive representation D may learn to prime the sensory representations S of all cues, including CSs, that have consistently been activated when it has; and S $\rightarrow$ M habit, or motor, learning whereby the sensory-motor maps, vectors, and gains that are involved in motor control may be adaptively calibrated.

These processes contribute to the modulation of declarative memory by motivational feedback and to the learning and performance of procedural memory. Thus learned S $\rightarrow$ D $\rightarrow$ S positive feedback quickly draws attention to motivationally salient cues and blocks activation of less salient cues via lateral inhibition among the S categories. D $\rightarrow$ S motivational feedback also energizes the release of discrete adaptive S $\rightarrow$ M responses. Based on a theoretical analysis, the final common path of the drive representations D, at or after the stage at which motivational decisions are made, was predicted to intersect or be modulated by the hippocampal formation (Grossberg, 1975, 1982). In support of this prediction, Thompson *et al.* (1984, 1987) have shown that emotional conditioning (as in the S $\rightarrow$ D circuit) influences hippocampal sites, whereas motor conditioning (as in the S $\rightarrow$ M circuit) occurs within the cerebellum. In addition, hippocampal ablation attenuates blocking (Rickert, Bennett, Lane, and French, 1978; Schmajuk, Spear, and Isaacson, 1983; Solomon, 1977). Blocking fails in the model when D $\rightarrow$ S feedback is impaired, as follows. In the complete model, when the S population activities that categorize conditioned reinforcers are amplified by strong conditioned S $\rightarrow$ D $\rightarrow$ S attentional feedback, they can block activation of other S populations via S $\rightarrow$ S lateral inhibition. When D $\rightarrow$ S feedback is removed, amplification and its blocking effect are eliminated. See Grossberg and Levine (1987) for blocking simulations. These model properties clarify how damage to the hippocampal system that involves both its drive-modulatory and orienting functions can result in either impaired or abnormally strong utilization of contextual cues (due to a failure of S $\rightarrow$ D $\rightarrow$ S feedback to enhance salient cues), and a failure of flexible reset and memory search for appropriate cues to attend (due to a failure of the orienting subsystem $\mathcal{A}$ to trigger these events).

Why should a single brain region, like the hippocampal system, modulate both recognition learning and reinforcement learning? We suggest that this is so in part because the same adaptive timing and orienting processes modulate both types of learning (Grossberg and Merrill, 1992; Grossberg and Schmajuk, 1989). This linkage clarifies how the hippocampal system may mediate tasks like delayed non-match to sample (DNMS) wherein both temporal delays and novelty-sensitive recognition processes are involved (Gaffan, 1974; Mishkin and Delacour, 1975). The proposed adaptive timing and orienting properties of the hippocampal system are envisaged to cooperate in the following way. As shown in Figures 3 and 4, $S \rightarrow D \rightarrow S$ feedback can rapidly focus attention on motivationally salient cues, as inhibition from D to the orienting subsystem inhibits orienting reactions that would otherwise occur in response to irrelevant situational cues. The inhibition from D to the orienting subsystem helps to model competition between consummatory and orienting behaviors (Staddon, 1983).

Another process is, however, needed to prevent the premature reset of attention by potentially distracting irrelevant cues during variable task-specific delays. For example, suppose that an animal inspects a food box right after a signal occurs that has regularly predicted food delivery in six seconds. When the animal inspects the food box, it perceives the nonoccurrence of food during the subsequent six seconds. These nonoccurrences disconfirm the animal's sensory expectation that food will appear in the magazine. Because the perceptual processing cycle that processes this sensory information occurs at a much faster rate than six seconds, it can compute this sensory disconfirmation many times before the six second delay has elapsed. Why is not the mismatch between the learned expectation of food and the percept of no-food treated like a predictive failure? Why, as often occurs when a previously rewarded cue is no longer rewarded, does the mismatch not trigger reset of attention, frustration, forgetting, and exploratory behavior? Were this to happen, humans and animals would restlessly explore their environments without being able to wait for delayed rewards.

The central issue is: What spares the animal from erroneously reacting to these *expected nonoccurrences* of food during the first six seconds as predictive failures? Why does the animal not immediately become so frustrated by the nonoccurrence of food that it shifts its attentional focus and releases exploratory behavior aimed at finding food somewhere else? Alternatively, if the animal does wait, but food does not appear after the six seconds have elapsed, why does the animal then react to the *unexpected nonoccurrence* of food by becoming frustrated, shifting its attention, and releasing exploratory behavior?

Grossberg and Schmajuk (1989) and Grossberg and Merrill (1992) argued that a primary role of the timing mechanism is to inhibit, or *gate*, the process whereby a disconfirmed expectation would otherwise negatively reinforce previous consummatory behavior, shift attention, and release exploratory behavior. The process of registering sensory mismatches or matches is not itself inhibited; if the food happened to appear earlier than expected, the animal could still perceive it and eat. Instead, the effects of these sensory mismatches upon reinforcement, attention, and exploration are inhibited.

Spectral Timing in the Hippocampus and Deficits due to its Removal

In order to realize this property, we suggested that a "spectral timing" circuit $S \rightarrow T$ operates in parallel with the fast $S \rightarrow D \rightarrow S$ emotional conditioning circuit (Figure 4) to maintain attention on salient cues during variable task-specific delays. Different populations of cells in T can be conditioned to respond selectively to different ISI intervals. The

total population output sums the output from all cells in the spectrum. Remarkably, this population response accurately models the ISI, even though no single cell does (Figure 5). Learned S → T timing *maintains* inhibition of the orienting subsystem and, in the example noted above, enables attention to be maintained on motivationally salient goal-related cues within the 6 second delay. If food does not occur even after 6 or more seconds have elapsed, then the adaptive timing circuit becomes quiet, and subsequent ART mismatches can trigger attentional reset, frustration, forgetting, and exploration in a manner modeled in Grossberg (1987).

We predicted in Grossberg and Merrill (1992) that this spectral timing circuit T exists in the hippocampal dentate-CA3 region in order to explain neurophysiological data showing that hippocampal CA3 pyramidal cell firing often mirrors the temporal delays observed in the conditioned nictitating membrane response (Berger, Berry, and Thompson, 1986). We suggested that subsets of hippocampal dentate cells respond at different rates to generate the spectral representation that controls the adaptively timed population response at CA3 pyramidal cells. Nowak and Berger (1992) have reported experimental evidence that is consistent with this prediction.

Figure 4

If the hippocampal system is removed, should animals and humans always have problems with DNMS and related tasks that involve stimulus delays? In the model, when the timing circuit T is removed, attention may more easily be distracted from goal objects during task-related delays. On the other hand, if the orienting subsystem is also removed, then flexible reset of attention in response to novel events is impaired, thereby eliminating a key mechanism whereby a distracting event could undermine performance. If the attentional system remains intact, then direct activation of individual recognition codes in response to a familiar event is still possible, and the matching process *per se* can partially update short term memory. However, the network can no longer flexibly search for the proper configuration of targets to attend, especially in the presence of complex spatial layouts that include distracting cues. The lack of timed control over variable delays can thus harm behavior more when it is necessary to shift attention among different sets of cues. Gaffan (1992) has described analogous data from hippocampectomized monkeys.

Both DNMS performance at brief delays and single-pair object discrimination learning with brief intertrial intervals are spared in hippocampal subjects (Eichenbaum, Otto, and Cohen, 1994). In the model, this is also true because the fast S → D → S attentional circuit remains intact. Long interstimulus delays, say of a day, also spare the performance of animals in some conditions (Mishkin, Malamut, and Bachevalier, 1984). These results have led some investigators to claim that the hippocampal system subserves a memory store of *intermediate* duration (Eichenbaum, Otto, and Cohen, 1994). As noted above, how the hippocampal system could create such a representation before it is transferred to the appropriate neocortical representations across several modalities has never been explained, and faces serious conceptual difficulties.

The ART model does not need to posit any such hippocampal memory store. At short delays, the fast feedback S → D → S system helps to focus attention on motivationally salient objects and to initiate attentional blocking. The failure of blocking at intermediate delays due to removal of the S → T circuit leads to abnormally strong utilization of contextual cues.

This processing failure causes little problem at long delays because potentially disruptive cues, being so widely separated in time, decay before they can compete for attention. These properties can be inferred from the model simulations of blocking by Grossberg and Levine (1987). It has not, to our knowledge, yet been tested whether the spectral timing circuit that is proposed to exist in dentate–CA3 plays the role described above in the DNMS paradigm.

Figure 5

Spectrally Timed Gain Control in the Cerebellum

Why is adaptive timing also needed in the motor conditioning circuit? This need is clarified by the fact that the S → D → S circuit focuses attention quickly on motivationally salient cues and can thereby just as quickly activate the motor circuit (Figure 3). Without adaptive timing within the motor circuit itself, the conditioned response could be prematurely released. Thus the clear survival advantage of attending quickly to motivationally important sensory events could disrupt the properly timed execution of responses contingent upon these events. The model suggests that this problem does not occur during normal behaviors because the hippocampal dentate-CA3 circuit and the cerebellar motor circuit are both adaptively timed. These distinct timing functions have been dissociated through ablation (Ebner and Bloedel, 1981; Gilbert and Thach, 1977; Optican and Robinson, 1980; Thompson, 1988; Thompson *et al.*, 1984, 1987) and ISI shift experiments during which the peak time of the hippocampal trace can change before the peak time of the discrete adaptive response (Hoehler and Thompson, 1980). The model suggests that orienting responses may be inhibited by the hippocampal dentate-CA3 timing circuit during the same time intervals when conditioned responses are disinhibited by the cerebellar timing circuit. This coordinated action extends the classical idea that consummatory and orienting responses are mutually inhibitory.

Recent experiments on conditioning the rabbit NMR suggest that response learning occurs within a subcortical cerebellar pathway, whereas response timing occurs within the cerebellar cortex (Perrett, Ruiz, and Mauk, 1993). If the cortical timing circuit is ablated, then motor responses are, indeed, prematurely released. These experimental results are consistent with the classical hypothesis that a fast cerebellar motor pathway—here interpreted to be subcortical (Lisberger, 1988)—can learn a conditioned gain appropriate to the response using climbing fiber inputs as a teaching signal (Albus, 1971; Fujita, 1982a, 1982b; Grossberg, 1969a, 1969b; Grossberg and Kuperstein, 1986; Marr, 1969).

We hypothesize, in addition, that adaptive timing is learned by a spectral timing circuit in which parallel fiber-Purkinje cell cortical synapses use climbing fiber inputs as a teaching signal (Figure 6). In this conception, cortical learning opens a timed gate by removing Purkinje cell inhibition from subcortical sites. As the timed gate opens, the subcortical motor pathway can read-out its learned gain with the correctly timed ISI between CS and US. Learned suppression of Purkinje cell output may be accomplished by conditioned long term depression, or LTD (Hoehler and Thompson, 1980; Ito, 1984). Eight key data properties have been simulated by this model (Bullock, Fiala, and Grossberg, 1994): Model Purkinje cell activity decreases in the interval following the onset of the CS, model nuclear cell responses match CR topography, CR peak amplitude occurs at the US onset, a discrete CR peak shift occurs with a change in ISI between CS and US, mixed training at two different ISIs produces a double-peaked CR, peak CR acquisition and response rates depend unimodally on the ISI,

CR onset latency decreases during training, and maladaptively-timed small-amplitude CRs result from ablation of cerebellar cortex.

Figure 6

Some striking cellular and circuit homologs exist between these model cerebellum and hippocampal timing mechanisms. Both control an inhibitory gate that modulates another learning process, and both occur on dendrites whose summed output across a spectrum of rate-sensitive cell sites determines the collective timed response. These similarities suggest the prediction that both the hippocampal dentate cell and cerebellar Purkinje cell dendrites may undergo similar biophysical events during conditioning.

Cooperative Hippocampal and Cerebellar Timing During Serial Compound Conditioning

How do the hippocampal and cerebellar timing circuits cooperate during timed behaviors? We illustrate such cooperation below by explaining paradoxical data about serial compound conditioning, during which a sequence CS_1–CS_2–US of two CS's precedes a US (Kehoe and Morrow; 1984; Kehoe *et al.*, 1979, 1987). Robust serial compound conditioning to CS_1 can occur even if primary CS_1–US conditioning at the same ISI, in the absence of CS_2, is ineffective. This happens, for example, if the CS_1–CS_2 ISI = 2400 msec and the CS_2–US ISI = 400 msec (Kehoe and Morrow, 1984). How does the occurrence of CS_2 enable CS_1 to bridge the 2800 msec ISI before US occurs?

We suggest that CS_2 can reactivate the sensory representation S_1 of CS_1 via the drive representation D along the feedback pathway $CS_1 \rightarrow S_1 \rightarrow D \rightarrow S_2$, and thereby restart the $S_1 \rightarrow T$ and $S_1 \rightarrow M$ timing circuits. In particular, on the first learning trial, the activity of S_1 does not persist until US occurs, but the activity of S_2 does. As a result, $S_2 \rightarrow D$ and $D \rightarrow S_2$ conditioning start to occur. On later learning trials, S_1 is active when CS_2 occurs. Thus S_1 is active when S_2 activates D. S_1 can hereby also learn to activate D, and D can be reciprocally conditioned to both S_1 and S_2 via the $D \rightarrow S_1$ and $D \rightarrow S_2$ feedback pathways. In this way, activation of D by CS_2 reactivates S_1 and restarts its timing circuits, so that they are active when the US occurs. As a result, $S_1 \rightarrow M$ conditioning of the NMR is possible, but is released earlier than the 2800 msec ISI between CS_1 and US.

This explanation clarifies why, if the ISI between CS_1 and CS_2 is short enough, then CS_2 elicits less NMR conditioning than it does when it is conditioned to the US at the same ISI without the occurrence of CS_1 (Kehoe *et al.*, 1979). If the CS_1–CS_2 delay is short enough, S_1 can partially block S_2 because $S_1 \rightarrow D \rightarrow S_1$ feedback is still strong when CS_2 occurs. Conversely, if the total CS_1–US ISI is increased, then CS_2 can elicit more NMR conditioning than it would in the absence of CS_1. Here, S_1's activity subsides by the time S_2 occurs, but it primes D with residual activity that can amplify $S_2 \rightarrow D \rightarrow S_2$ and $S_2 \rightarrow T$ conditioning when CS_2 and US occur. Kehoe *et al.* (1993) have shown that a spectral timing model can, indeed, be used to simulate key properties of serial compound conditioning data.

PART II

START: A Unified Model of Adaptive Timing and Conditioned Reinforcer Learning

The hippocampal adaptive timing model depicted in Figure 4 will now be mathematically defined. It combines Spectral Timing mechanisms with mechanisms from Adaptive Resonance Theory. Hence it is called the START model (Grossberg and Merrill, 1992). The START model builds upon a previous model of reinforcement learning whose processing stages have been compared with behavioral and neural data in a series of previous articles. Here we provide just enough review and exposition to define the model and to compare its emergent properties with illustrative data.

Although the model is helpful for the explanation of both classical and operant conditioning data, here each conditionable sensory event is called a conditioned stimulus, or CS. The i^{th} sensory event is denoted by CS_i. Event CS_i activates a population of cells that is called the i^{th} sensory representation S_i (Figure 4). Another population of cells, called a drive representation D, receives a combination of sensory, reinforcement, and homeostatic (or drive) stimuli. Reinforcement learning, emotional reactions, and motivational decisions are controlled by D (Grossberg, 1971, 1972a, 1982b). In particular, a reinforcing event, such as an unconditioned stimulus, or US, is capable of activating D.

Various authors have invoked representations analogous to drive representations. Bower and his colleagues have called them emotion nodes (Bower, 1981; Bower, Gilligan, and Monteiro, 1981) and Barto, Sutton, and Anderson (1983) have called them adaptive critic elements. During conditioning, presentation of a CS_i before a US causes activation of S_i followed by activation of D. Such pairing causes strengthening of the adaptive weight, or long term memory trace, in the modifiable synapses from S_i to D. This learning event converts CS_i into a conditioned reinforcer. Conditioned reinforcers hereby acquire the power to activate D via the conditioning process.

In the START model, reinforcement learning in $S_i \rightarrow D$ pathways is supplemented by a parallel learning process that is concerned with adaptive timing. As shown in Figure 4, both of these learning processes output to D, which in turn inhibits the population of cells in the orienting subsystem $\mathcal{A}$. The orienting subsystem is a source of nonspecific arousal signals that are capable of initiating frustrative emotional reactions, attention shifts, and orienting responses. The inhibitory pathway from D to $\mathcal{A}$ is the gate that prevents these events from occurring.

Limited Capacity Short Term Memory

The sensory representations S_i compete for a limited capacity, or finite total amount, of activation. Winning populations are said to be stored in short term memory, or STM. The competition is carried out by an on-center off-surround interaction among the populations S_i. The property of STM storage is achieved by using recurrent, or feedback, pathways among the populations. A tendency to select winning populations is achieved by using membrane equations, or shunting interactions, to define each population's activation, and a proper choice of feedback signals between populations (Grossberg, 1973, 1982b). Expressed mathematically, each CS_i activates an STM representation S_i whose activity S_i obeys the

shunting on-center off-surround competitive feedback equation

$$\frac{d}{dt}S_i = -\alpha_A S_i + \beta_A(1 - S_i)(I_i(t) + f_S(S_i)) - \gamma_A S_i \sum_{k \neq i} f_S(S_k). \tag{2}$$

In (2), $I_i(t)$ is the input that is turned on by presentation of CS_i. Term $-\alpha_A S_i$ describes passive decay of activity S_i. Term $\beta_A(1 - S_i)(I_i(t) + f_S(S_i))$ describes the excitatory effect on S_i of the input $I_i(t)$ and the feedback signal $f_S(S_i)$ from population S_i to itself. Activity S_i can continue to grow until it reaches the excitatory saturation point, which is scaled to equal 1 in (2). Term $-\delta_A S_i \sum_{k \neq i} f_S(S_k)$ describes inhibition of S_i by competitive signals $f_S(S_k)$ from the off-surround of populations $k \neq i$. Figure 7 summarizes a computer simulation of how a brief CS_1 gives rise to a sustained STM activation S_1, which is partially inhibited by competition from S_0's activation in response to a US. The signal function f_S may be chosen to have any of the several forms without qualitatively altering model properties. In this chapter, the simple threshold-linear half-wave rectification function

$$f(w) \quad [w - \mu]^+ \equiv \max(w - \mu, 0) \tag{3}$$

is used, except in equation (9) below, which uses a sigmoid, or S-shaped, signal function.

Figure 7

Drive Representation

The computer simulations reported herein use only a single drive representation D. Explanations of data arising from competing drive representations are discussed in Grossberg (1984, 1987). The activity D of the drive representation D obeys the equation

$$\frac{d}{dt}D = -\alpha_D D + \beta_D \sum_i f_D(S_i)C_i + \gamma_D R. \tag{4}$$

In (4), term $-\alpha_D D$ describes the passive decay of activity D. Term $\beta_D \sum_i f_D(S_i)C_i$ describes the total excitatory effect of all the sensory representations S_i on D. In this term, the signal function f_D is chosen as in (3), and C_i is the adaptive weight, or long term memory (LTM) trace, in the pathway from the sensory representation S_i of CS_i to the drive representation D. This LTM trace is denoted by C_i because its size measures how well S_i can activate D, and thus how CS_i ($i \geq 1$) has become a conditioned reinforcer through learning. Because C_i multiplies $f_D(S_i)$, a large activation of S_i will have a negligible effect on D if C_i is small, and a large effect on D if C_i is large. Coefficient C_0 is set equal to a large value from the start because it enables the US to activate D via its sensory representation S_0. Term $\gamma_D R$ describes the total output of the spectral timing circuit to D. Output R is defined in (12).

Figure 8c summarizes a computer simulation in which the activity D responds to CS and US signals after 50 conditioning trials. Figures 8a and 8b summarize the corresponding STM traces S_1 of the CS and S_0 of the US, respectively.

Figure 8

Conditioned Reinforcement

The adaptive weight C_i that calibrates conditioned reinforcement obeys a gated learning law (Grossberg, 1969b):

$$\frac{d}{dt}C_i = \alpha_C S_i(-C_i + \beta_C(1 - C_i)f_C(D)). \tag{5}$$

Learning by C_i is turned on and off by the signal S_i from S_i, which thus acts like a learning gate, or modulator. Once turned on, C_i performs a time-average of activity at the drive representation D via the signal $f_C(D)$, which is chosen as in (3). Activity C_1 cannot exceed the finite value 1, due to the shunting term $1 - C_i$. The value of C_i can both increase and decrease during the course of learning. The remaining equations of the model describe the adaptive timing process.

Activation Spectrum

The START model is said to control "spectral" timing because each drive representation D is associated with a population of cell sites whose members react at a spectrum of rates r_j. Neural populations whose elements are distributed along a temporal or spatial parameter are familiar throughout the nervous system. Two examples are populations of spinal cord cells that obey the size principle (Henneman, 1957, 1985), and the spatial frequency-tuned cells of the visual cortex (Jones and Keck, 1978; Musselwhite and Jeffreys, 1985; Parker and Salzen, 1977a, 1977b; Parker, Salzen, and Lishman, 1982a, 1982b; Plant, Zummern, and Durden, 1983; Skrandies, 1984; Vassilev, Manahilov, and Mitov, 1983; Vassilev and Strashimirov, 1979; Williamson, Kaufman, and Brenner, 1978).

The spectral activities x_{ij} that are associated with drive representation D and activated by sensory representation S_i obey the equation

$$\frac{d}{dt}x_{ij} = r_j(-x_{ij} + (1 - x_{ij})f_x(S_i)), \tag{6}$$

where f_x satisfies (3). By (2) and (6), presentation of CS_i to S_i via an input I_i generates an output signal $f_x(S_i)$ that activates the local potentials x_{ij} of all cell sites in the target population. The potentials x_{ij} respond at rates proportional to r_j, $j = 1, 2, \ldots, n$. These potentials activate the next processing stage via signals

$$f(x_{ij}) = \frac{x_{ij}^8}{\delta_{ij}^8 + x_{ij}^8}. \tag{7}$$

Signal $f(x_{ij})$ is a sigmoid function of activity x_{ij}. Figure 5a shows the activation spectrum $f(x_{ij}(t))$ that arises from presentation of CS_i to S_i via input I_i in (2), using a choice of rate parameters r_j in (6) which range from 10 (fast) to 0.0025 (slow).

Habituative Transmitter Spectrum

Each spectral activation signal $f(x_{ij})$ interacts with a habituative chemical transmitter y_{ij} via the equation

$$\frac{d}{dt}y_{ij} = \alpha_y(1 - y_{ij}) - \beta_y f(x_{ij})y_{ij}. \tag{8}$$

According to equation (8), the amount of neurotransmitter y_{ij} accumulates to a constant target level 1, via term $\alpha_y(1 - y_{ij})$, and is inactivated, or *habituates*, due to a mass action

interaction with signal $f(x_{ij})$, via term $-\beta_y f(x_{ij})y_{ij}$. The different rates r_j at which each x_{ij} is activated causes the corresponding y_{ij} to become habituated at different rates. The family of curves $y_{ij}(t)$, $j = 1, 2, \ldots, n$, is called a habituation spectrum. The signal functions $f(x_{ij}(t))$ in Figure 5a generate the habituation spectrum of $y_{ij}(t)$ curves in Figure 5b.

Gated Signal Spectrum

Each signal $f(x_{ij})$ interacts with y_{ij} via mass action. This process is also called *gating* of $f(x_{ij})$ by y_{ij} to yield a net signal g_{ij} that is equal to $f(x_{ij})y_{ij}$. Each gated signal $g_{ij}(t) \equiv f(x_{ij}(t))y_{ij}(t)$ has a different rate of growth and decay, thereby generating the gated signal spectrum shown in Figure 5c. In these curves, each function $g_{ij}(t)$ is a unimodal function of time, where function $g_{ij}(t)$ achieves its maximum value M_{ij} at time T_{ij}, T_{ij} is an increasing function of i, and M_{ij} is a decreasing function of j.

These laws for the dynamics of a chemical transmitter were described in Grossberg (1968b, 1969a). They capture the simplest first-order properties of a number of known transmitter regulating steps (Cooper, Bloom, and Roth, 1974), such as transmitter production (term α_y), feedback inhibition by an intermediate or final stage of production on a former stage (term $-\alpha_y y_j$), and mass action transmitter inactivation (term $-\beta_y f(x_j)y_j$). Alternatively, they can be described as the voltage drop across an RC circuit, or the current flow through an appropriately constructed transistor circuit. These properties are sufficient to explain the article's targeted data, so finer transmitter processes, such as transmitter mobilization effects, are not considered herein.

Spectral Learning Law

Learning of spectral timing obeys a gated steepest descent equation

$$\frac{d}{dt}z_{ij} \quad \alpha_z f(x_{ij})y_{ij}(-z_{ij} + N), \tag{9}$$

where N is the Now Print signal that is below defined in (10). Each *long term memory (LTM) trace* z_{ij} in (9) is activated by its own sampling signal $g_{ij} = f(x_{ij})y_{ij}$. The sampling signal g_{ij} turns on, or *gates*, the learning process, and causes z_{ij} to approach N during the sampling interval at a rate proportional to g_{ij}. The attraction of z_{ij} to N is called *steepest descent*. Thus (9) is a variant of the gated steepest descent equation that was defined in (1). Each z_{ij} changes by an amount that reflects the degree to which the curves $g_{ij}(t)$ and $N(t)$ have simultaneously large values through time. If g_{ij} is large when N is large, then z_{ij} increases in size. If g_{ij} is large when N is small, then z_{ij} decreases in size. As in equation (5), z_{ij} can either increase or decrease as a result of learning.

Now Print Signal

A transiently active Now Print signal N modulates the learning process of (9). The signal N may be activated either by a US or by a CS that has already become a conditioned reinforcer. Figure 4 shows that both the US and a conditioned reinforcer CS can activate the drive representation D. Equation (4) describes this property mathematically. The Now Print signal N (for example, in Figure 9c) is turned on by sufficiently large and rapid increments in the activity D of D (for example, in Figure 9b). As Figures 9b and 9c illustrate, the Now Print signal N approximates the time derivative of the drive activity D. A neurophysiologically plausible way to achieve this property is to assume that the transient signal N is derived from the sustained activity D by the action of a slow inhibitory interneuron (see

Figure 9a). The transformation from sustained activity D to transient activity N can be realized mathematically by the function

$$N = [f_C(D) - E - \epsilon]^+.$$

In (10), E is the activity of an inhibitory interneuron that time-averages $f_C(D)$, as in equation

$$\frac{d}{dt}E = \alpha_E(-E + f_C(D)), \tag{11}$$

before inhibiting the direct excitatory signal $f_C(D)$. Equation (10) means that $N = 0$ if $f_C(D) - E \leq \epsilon$, and $N = f_C(D) - E - \epsilon$ if $f_C(D) - E > \epsilon$. By (10), N responds rapidly to an increment in $f_C(D)$. By (11), the inhibitory interneuron activity E responds more slowly to $f_C(D)$. As E grows, it inhibits the influence of $f_C(D)$ on N, by (10), thereby shutting N off. As noted below, an important property of N is that it increases in amplitude, but not significantly in duration, in response to larger inputs $f_C(D)$.

Figure 9

As noted above, the time interval between CS onset and US onset is called the *interstimulus interval*, or ISI. Using the spectral learning law (9)–(11), the individual LTM traces differ in their ability to learn at different values of the ISI. This is the basis of the network's timing properties. Figure 10 illustrates how six different LTM traces z_j, $i = 1, \ldots, 6$, learn during this simulated learning experiment. The CS and US are paired during 4 learning trials, after which the CS is presented alone on a single performance trial. In this computer simulation, the CS input $I_{CS}(t)$ remained on for a duration of 0.05 time units on each learning trial. The US input $I_{US}(t)$ was presented after an ISI of 0.5 time units and remained on for 0.05 time units. The upper panel in each part of the figure depicts the gated signal function $g_{ij}(t)$ with r_j chosen at progressively slower rates. The middle panel plots the corresponding LTM trace $z_{ij}(t)$.

Figure 10

Doubly Gated Signal Spectrum

The lower panel plots the twice-gated signal $h_{ij}(t) = f(x_{ij}(t))y_{ij}(t)z_{ij}(t)$. Each twice-gated signal function $h_{ij}(t)$ registers how well the timing of CS and US is learned and read-out by the i^{th} processing channel. In Figure 10b, where the once-gated signal $g_{ij}(t)$ peaks at approximately the ISI of 0.5 time units, the LTM trace $z_{ij}(t)$ shows the maximum learning. The twice-gated signal $h_{ij}(t)$ also shows a maximal enhancement due to learning, and exhibits a peak of activation at approximately 0.5 time units after onset of the CS on each trial. This behavior is also generated on the fifth trial, during which only the CS is presented.

J. Output Signal

The output of the network is the sum of the twice-gated signals $h_{ij}(t)$ from all the spectral components corresponding to all the CS_i. Thus

$$R = \sum_{i,j} f(x_{ij})y_{ij}z_{ij}$$

The output signal computes the cumulative learned reaction of the whole population to the input pattern. Figure 5e shows the function R derived from the pooled signals h_{ij} shown in Figure 5d. A comparison of Figures 5c-e illustrate how the output $R(t)$ generates an accurately timed response from the cumulative partial learning of all the cell sites in the population spectrum. The once-gated signals $g_{ij}(t)$ in Figure 5c are biased towards early times. The twice-gated signals $h_{ij}(t)$ in Figure 5d are biased towards the ISI, but many signals peak at other times. The output $R(t)$ in Figure 5e combines these partial views into a cumulative response that peaks at the ISI.

The Problem of Self-Printing During Adaptively Timed Secondary Conditioning

The START model of Grossberg and Merrill (1992) overcame four types of problems whose solution is needed to explain behavioral and neural data about adaptively timed conditioning. These are the problems of (1) self-printing during adaptively timed secondary conditioning, (2) asymmetric effects of increasing CS or US intensity on timed responding, (3) different effects of US duration on timing than on reinforcement, and (4) combinatorial explosion of network pathways. These problems and their solution by the START model are reviewed below, along with supportive data. Problems (1), (3), and (4) were not solved by the Grossberg and Schmajuk (1989) model.

A major problem for any model of adaptive timing is to explain how adaptively timed secondary conditioning can occur. In primary conditioning, a conditioned stimulus CS_1 is paired with an unconditioned stimulus US until CS_1 becomes a conditioned reinforcer. In secondary conditioning, another conditioned stimulus CS_2 is paired with CS_1 until it, too, gains reinforcing properties. Various experiments have shown that the conditioned response to CS_2 can be adaptively timed (Gormezano and Kehoe, 1984; Kehoe, Marshall-Goodell, and Gormezano, 1987). Indeed, Gormezano and Kehoe (1984) claimed that, in their experimental paradigm, "first- and second-order conditioning follow the same laws" (p. 314), although they also acknowledged that some variables may differentially effect first-order and second-order conditioning in other paradigms.

Adaptively timed secondary conditioning could easily erase the effects of adaptively timed primary conditioning in the following way. In order for CS_1 to act as a conditioned reinforcer, CS_1 must gain control of the pathway along which the US activates its reinforcing properties. Suppose that CS_1 activated its sensory representation S_1 via an input (I_{CS_1}) pathway and that US expressed its reinforcing properties via an input (I_{US}) pathway. Also suppose that conditioned reinforcer learning enabled CS_1 to activate I_{US}. Thereafter, presentation of CS_1 would *simultaneously* activate both the I_{CS_1} pathway and the I_{US} pathway. This coactivation would create new learning trials for CS_1 with a *zero* ISI. In other words CS_1 could *self-print* a spectrum with zero ISI due to CS_1–CS_1 pairing via the I_{CS} and conditioned I_{US} pathway. Thus, as CS_1 became a conditioned reinforcer, it could undermine the timing that it learned through CS_1–US pairing during primary conditioning. Such self-printing could, for example, occur on secondary conditioning trials when a CS_2 is followed by a conditioned reinforcer CS_1.

Simulations of Secondary Conditioning

The START model overcomes the self-printing problem with its use of a transient Now Print signal N, as in (10). During primary conditioning, onset of the US causes a brief

output burst from N. During secondary conditioning, onset of the conditioned reinforcer CS_1 also causes a brief output burst from N. However, the spectrum activated by CS_1 takes awhile to build up, so essentially all of its activities x_{ij} and sampling signals $f(x_{ij})y_{ij}$ are very small during the brief interval when N is large (compare Figures 5c and 9c). By the spectral learning law (9), negligible self-printing occurs. The main effect of the self-printing that does occur is to reduce every spectral LTM trace z_{1J} in (9) by a fixed proportion of its value, thus scaling down the size of $R(t)$ without changing the timing of its peak.

Figure 1

Figure 11a depicts the model output $R(t)$ when the Now Print threshold ϵ in (10) is set to a high enough level to guarantee that no self-printing or secondary conditioning occur. Here CS_1 never activates a Now Print signal. Figure 11b shows the output when ϵ is set lower, thus allowing secondary conditioning and some self-printing to occur. Correct timing still obtains.

Figure 12

Figure 12 shows how the model behaves during secondary conditioning. The left hand half of each panel shows the output of the model in response to the primary conditioned stimulus CS_1, and the right hand half of each panel shows the model output in response to the secondary conditioned stimulus CS_2. The peak time arising from the presentation of CS_2 occurs near the expected time of arrival of CS_1, rather than the expected time of the US. This property is consistent with the environment that a model or animal experiences, since the subject never sees CS_2 paired with the primal US, but rather sees it paired as a predictor of CS_1, which serves as a CR in this context.

Asymmetric CS and US Processing in Timing Control

Although CS_1 can attain properties of a conditioned reinforcer through CS_1–US pairing, this does not imply that all the functional properties of a conditioned reinforcer and an unconditioned stimulus are interchangeable. For example, increasing the intensity of a conditioned reinforcer CS_1 can "speed up the clock" (Maricq, Roberts, and Church, 1981; Meck and Church, 1987; Wilkie, 1987), whereas increasing the intensity of a primary US can increase the amplitude of conditioned response, but does not change its timing (Smith, 1968).

The way that parametric changes of CS and US cause different effects on adaptive timing places constraints on possible mechanisms of how adaptive timing is learned during secondary conditioning. Although the CS acquires reinforcing properties of a US when it becomes a conditioned reinforcer, it may not acquire all of its timing properties. The proposed solution of the self-printing problem in Grossberg and Merrill (1992) suggests how different responses may be caused by an increase in CS intensity or US intensity. This explanation holds even if the CS_1 and US sensory representations S_1 and S_0, respectively, each send signals along the same types of pathways to the drive representation and the adaptive timing circuit. The explanation is summarized below.

An increase in CS_1 intensity causes an increase in the amplitude of input $I_1(t)$ in (2). The larger input causes a larger peak amplitude of activity S_1 in (2), and a larger signal $f_x(S_1)$ in (6). By (6), the rate with which a spectral activation x_{1j} reacts to signal $f_x(S_1)$

equals $r_j(1 + f_x(S_1))$. Thus an increase in CS_1 intensity speeds up the processing of *all* spectral activations x_{1j}. Because CS_1 is a conditioned reinforcer, some of its LTM traces z_{1j} are nonzero. Thus the total output R in (12) peaks at an earlier time, and causes the total output D from D in (4) to also peak at an earlier time.

In contrast, a primary reinforcer such as a US does not generate a significant output $R(t)$ from the spectral timing circuit, even if it is allowed to generate a large signal $f_x(S_0)$ to the adaptive timing circuit in (6). This is true because a large US generates a signal $f_x(S_0)$ to the spectral activations in (6) at the same time that it generates a large signal $f_D(S_0)$ to D in (4) and a large Now Print signal N in (9). Thus a US creates the conditions of a "zero ISI experiment" for purposes of spectral learning. All the LTM traces z_{0j} in (9) therefore remain very small in response to any number of US representations. An increase in US amplitude thus cannot cause speed-up of the output $R(t)$ in (12), because this output remains approximately zero in response to any US intensity. In summary, the same mechanism that explains how the self-printing problem is avoided also explains why an increase in CS intensity, but not US intensity, speeds up the conditioned response.

The primary effect of an increase in US intensity is to increase the amplitude of the signal $f_D(S_0)$ in (4) to the drive representation D. This causes an increase in the amplitude of D and thus an increase in the amplitude of the conditioned response that is modulated by D. This explanation of how a US increases the amplitude of the conditioned response also holds if the US sends no signal $f_x(S_0)$ directly to the adaptive timing circuit. Grossberg and Schmajuk (1989) further discuss this issue.

Different Effects of US Duration on Emotional Conditioning and Adaptive Timing: Sustained and Transient Responses

The existence of a transient Now Print signal N plays a central role in our explanations of how to avoid self-printing during secondary conditioning, and of different effects of CS and US intensity on learned timing. Another type of data lends support to the hypothesis that the activity D and the Now Print signal N both exist but respond to the US in different ways. These data show that an increase in US duration can significantly increase the strength of emotional conditioning (Ashton, Bitgood, and Moor, 1969; Boe, 1966; Borozci, Storms, and Broen, 1964; Church, Raymond, and Beauchamp, 1967; Keehn, 1963; Strouthes, 1965). How can a brief Now Print signal N whose duration does not increase significantly with US duration coexist with emotional conditioning properties that do increase significantly with US duration?

An answer can be given using properties of drive representations D. The activation D of a drive representation by a US does persist longer when the US duration is increased, and does thereby increase the strength of emotional conditioning at the S $\rightarrow$ D synapses that are modelled by equations (4) and (5); see Grossberg (1972a, Section 4 and Grossberg (1982a) for further discussion of this property. This sustained activation D of a drive representation gives rise to a transient Now Print signal N at a different processing stage—a transient detector—that is downstream from D itself, as displayed in Figures 4 and 9. Thus D and N represent responses of "sustained cells" and "transient cells"—a distinction familiar from visual perception—which here instantiate different functional properties of emotional conditioning and conditioning of adaptive timing, respectively. The parametric data properties summarized above illustrate that the processes of emotional conditioning and adaptive

timing, although linked, are not the same.

The Problem of Combinatorial Explosion: Stimulus versus Drive Spectra

According to any spectral timing theory, each CS_i activates a sensory representation S_i that broadcasts signals along many parallel pathways. This can lead to a combinatorial explosion of cell bodies if the spectra are incorrectly instantiated. For example, suppose that each pathway activated a different cell, and that each cell's activity computed a different x_{ij}, $j = 1, 2, \ldots, n$. Then there would exist as many copies of the spectral timing model as there are sensory representations in the brain. In addition, each spectrum contains 80 activities x_{ij} in our computer simulations. Such a model would require a huge number of cells to represent a different spectrum for every possible sensory representation. This is, in fact, the type of circuit used in the Grossberg-Schmajuk model.

In the START model, each *drive* representation, not every *sensory* representation, has its own spectral cells. Thus the pathways from all sensory representations that correspond to any given drive representation share the same neurons. This modification greatly reduces the number of cells that are needed to achieve spectral timing of arbitrary conditionable CS–US combinations, since there are many fewer drive representations (e.g., for hunger, thirst, sex, etc.) than there are sensory representations. As in Figure 4, each spectrum is computed in parallel with its drive representation. Since the present simulations consider only one type of reinforcer, only one drive representation is depicted. In general, each CS sends an adaptive pathway to every drive representation to which it can be conditioned, as well as adaptive pathways sufficient to sample the corresponding spectral representation. The "coordinates" of each drive representation and its spectrum encode reinforcement and homeostatic variables. In contrast, the CS-activated pathways to these circuits carry signals that reflect the sensory features of the CSs.

Thus the fact that different perceptual stimuli may elicit characteristic responses at the cells which represent adaptive timing does not, in itself, imply that these perceptual stimuli are "encoded" at those cells. Grossberg and Merrill (1992) suggest how hippocampal cells can form an adaptive timing circuit, and how dendrites of hippocampal pyramidal cells can represent a drive-based spectrum that avoids the combinatorial explosion in the dentate-CA3 circuit. This interpretation is used to suggest an explanation of anatomical, neurophysiological, and neuropharmacological data about this hippocampal circuit that are consistent with a functional role in adaptive timing. The Grossberg and Merrill (1992) article also provides computer simulations of data from experiments employing partial reinforcement (Roberts, 1981), ISI shifts (Coleman and Gormezano, 1971), time-averaging due to multiple CSs (Holder and Roberts, 1985), and multiple ISIs (Millenson, Kehoe, and Gormezano, 1977; Smith, 1968).

Concluding Remarks

The neural model described herein suggests how the hippocampal system and cerebellum may cooperate to control adaptively timed recognition learning, motivated attention, and conditioned responding. The model clarifies how the hippocampal system may combine novelty-based modulation of recognition learning and reinforcement learning with a competence for adaptively timed attention and inhibition of orienting responses. In particular, it suggests how orienting responses may be inhibited by the hippocampal dentate-CA3 timing

circuit during the same time interval during which goal-oriented conditioned responses are released by adaptively timed opening of the cerebellar Purkinje cell gate.

The model distinguishes between the micro-timing that is needed to determine how long motivated attention needs to be focused on a single predicted goal event, and the macro-timing whereby attention is maintained during the planned performance of a sequence of actions leading to a goal. Both sorts of timing would appear to be at work during many behaviors. A partially developed theory of how they are coordinated clarifies some aspects of the complex pattern of connections that exists between the temporal cortex, frontal cortex, and hippocampal system.

Why the hippocampal system should play a role in spatial orientation is also consistent with this modelling framework. This link is established when one poses the question of how an animal can direct its goal-oriented attentive behaviors among sets of environmental landmarks that vary in their motivational salience. Such a perspective is consistent with the proposal that the hippocampal system can play a role as a cognitive map (Leonard and McNaughton, 1990; O'Keefe, 1990; O'Keefe and Nadel, 1978), suitably defined, without denying its relevance for the control of approach-avoidance behaviors (Amsel, 1993). How to computationally integrate the steering role of reinforcement and motivation into a self-organizing network for spatial orientation remains an open problem. Despite these theoretical gaps, the ART models that have already been developed put mechanistic flesh on the metaphorical bones of declarative memory and procedural memory by articulating new behavioral principles, neural mechanisms, and experimental explanations and predictions that can be used to clarify how a freely moving individual flexibly learns about and acts upon valued goal objects in a timely fashion.

REFERENCES

Aggleton, J.P. (1993). The contribution of the amygdala to normal and abnormal emotional states. *Trends in Neurosciences*, **16**, 328–333.

Albus, J.S. (1971). A theory of cerebellar function. *Mathematical Biosciences*, **10**, 25–61.

Amsel, A. Hippocampal function in the rat: Cognitive mapping or vicarious trial and error. (1993) *Hippocampus*, **3**, 251–256.

Artola, A. and Singer, W. (1993). Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in Neurosciences*, **16**, 480–487.

Asfour, Y.R. (1994). Neural networks for multi-sensor fusion and classification. PhD Thesis, Department of Cognitive and Neural Systems, Boston University.

Asfour, Y.R., Carpenter, G.A., and Grossberg, S. (1995). Landsat image segmentation using the fuzzy ARTMAP neural network. **Technical Report CAS/CNS-TR-95-004**, Boston, MA: Boston University. In **Proceedings of the world congress on neural networks**, Washington, in press.

Asfour, Y.R., Carpenter, G.A., Grossberg, S., and Lesher, G. (1993). Fusion ARTMAP: A neural network architecture for multi-channel data fusion and classification. **Proceedings of the world congress on neural networks, II**, 210–215. Hillsdale, NJ: Erlbaum Associates.

Ashton, A.B., Bitgood, S.C., and Moore, J.W. (1969). Auditory differential conditioning of the rabbit nictitating membrane response: III. Effects of US shock intensity and duration. *Psychonomic Science*, **15**, 127–128.

Atkinson, R.C. and Shiffrin, R.M. (1971). The control of short term memory. *Scientific American*, August, 82–90.

Bachelder, I.A., Waxman, A.M., and Seibert, M. (1993). A neural system for mobile robot visual place learning and recognition. In **Proceedings of the world congress on neural networks, I**, 512–517. Hillsdale, NJ: Erlbaum Associates.

Baddeley, A.D. (1986). **Working memory**. Oxford: Clarendon Press.

Baloch, A.A. and Waxman, A.M. (1991). Visual learning, adaptive expectations, and behavioral conditioning of the mobile robot MAVIN. *Neural Networks*, **4**, 271–302.

Barto, A.G., Sutton, R.S., and Anderson, C.W. (1983). Neuron-like adaptive elements that can solve difficult learning problems. *IEEE Transactions*, **SMC-13**, 834–846.

Berger, T.W., Berry, S.D., and Thompson, R.F. (1986). Role of the hippocampus in classical conditioning of aversive and appetitive behaviors. In **The hippocampus, Volume 4** (R.L. Isaacson and K.H. Pribram, Eds). New York: Plenum Press, pp. 203–239.

Boardman, I. and Bullock, D. (1991). A neural network model of serial order recall from short-term memory. In **Proceedings of the international joint conference on neural networks**, Seattle, **II**, 879–884. Piscataway, NJ: IEEE Service Center.

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

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

Bower, G.H. (1981). Mood and memory. *American Psychologist*, **36**, 129–148.

Bower, G.H., Gilligan, S.G., and Monteiro, K.P. (1981). Selectivity of learning caused by adaptive states. *Journal of Experimental Psychology: General*, **110**, 451–473.

Bradski, G., Carpenter, G.A. and Grossberg, S. (1992). Working memory networks for learning temporal order with application to 3-D visual object recognition. *Neural Computation*, **4**, 270–286.

Bradski, G., Carpenter, G.A., and Grossberg, S. (1994). STORE working memory networks for storage and recall of arbitrary temporal sequences. *Biological Cybernetics*, **71**, 469–480.

Bradski, G. and Grossberg, S. (1994). A neural architecture for 3-D object recognition from multiple 2-D views. In **Proceedings of the world congress on neural networks, IV**, 211–219. Hillsdale, NJ: Erlbaum Associates.

Bradski, G. and Grossberg, S. (1995). Fast learning VIEWNET architectures for recognizing 3-D objects from multiple 2-D views. **Technical Report CAS/CNS-TR-93-053**, Boston, MA: Boston University. *Neural Networks*, in press.

Bruner, J.S. (1969). **The pathology of memory** (G.A. Talland and N.C. Waugh, Eds.). New York: Academic Press.

Bullock, D., Fiala, J.C., and Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Networks*, **7**, 1101–1114.

Carpenter, G.A. and Grossberg, S. (1987a). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, **37**, 54–115.

Carpenter, G.A. and Grossberg, S. (1987b). ART 2: Stable self-organization of pattern recognition codes for analog input patterns. *Applied Optics*, **26**, 4919–4930.

Carpenter, G.A. and Grossberg, S. (1990). ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, **3**, 129–152.

Carpenter, G.A. and Grossberg, S. (Eds.) (1991). **Pattern recognition by self-organizing neural networks**. Cambridge, MA: MIT Press.

Carpenter, G.A. and Grossberg, S. (1992). Fuzzy ARTMAP: Supervised learning, recognition, and prediction by a self-organizing neural network. *IEEE Communications Magazine*, **30**, 38–49.

Carpenter, G.A. and Grossberg, S. (1993). Normal and amnesic learning, recognition, and memory by a model of cortico-hippocampal interactions. *Trends in Neurosciences*, **16**, 131–137.

Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H., and Rosen, D.B. (1992). Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, **3**(5), 698–713.

Carpenter, G.A., Grossberg, S., and Reynolds, J.H. (1991). ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural Networks*, **4**, 565–588.

Carpenter, G.A., Grossberg, S., and Reynolds, J.H. (1995). A fuzzy ARTMAP nonparametric probability estimator for nonstationary pattern recognition. *IEEE Transactions on Neural Networks*, in press.

Carpenter, G.A. and Ross, W.D. (1993). ART-EMAP: A neural network architecture for learning and prediction by evidence accumulation. In **Proceedings of the world congress on neural networks, III**, 649–656. Hillsdale, NJ: Erlbaum Associates.

Carpenter, G.A. and Tan, A.-H. (1993). Rule extraction, fuzzy ARTMAP, and medical databases. In **Proceedings of the world congress on neural networks, I**, 501–506. Hillsdale, NJ: Erlbaum Associates.

Caudell, T.P., Smith, S.D.G., Escobedo, R., and Anderson, M. (1994). NIRS: Large-scale ART 1 neural architectures for engineering design retrieval. *Neural Networks*, **7**, 1339–1350.

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

Cohen, M.A. and Grossberg, S. (1986). Neural dynamics of speech and language coding: Developmental programs, perceptual grouping, and competition for short term memory.

Human Neurobiology, **5**, 1–22.

Cohen, M.A. and Grossberg, S. (1987). Masking fields: A massively parallel neural architecture for learning, recognizing, and predicting multiple groupings of patterned data. *Applied Optics*, **26**, 1866–1891.

Cohen, N.J. (1984). Preserved learning capacity in amnesia: Evidence for multiple memory systems. In **The neuropsychology of memory**, (L. Squire and N. Butters, Eds.). New York: Guilford Press, pp. 83–103.

Cohen, N.J. and Squire, L.R. (1980). Preserved learning and retention of a pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, **210**, 207–210.

Coleman, S.R. and Gormezano, I. (1971). Classical conditioning of the rabbit's (*Oryctocagus cuniculus*) nictitating membrane response under symmetrical CS-US interval shifts. *Journal of Comparative and Physiological Psychology*, **77**, 447–455.

Cooper, J.R., Bloom, F.E., and Roth, R.H. (1974). **The biochemical basis of neuropharmacology (second edition)**. New York: Oxford University Press.

Damasio, A.R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, **1**, 123–132.

Deadwyler, S.A., West, M.O., and Lynch, G. (1979). Activity of dentate granule cells during learning: Differentiation of perforant path inputs. *Brain Research*, **169**, 29–43;

Deadwyler, S.A., West, M.O., and Robinson, J.H. (1981). Entorhinal and septal inputs differentially control sensory-evoked responses in the rat dentate gyrus. *Science*, **211**, 1181–1183.

Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience*, **3**, 1–8.

Desimone, R. and Ungerleider, L.G. (1989). Neural mechanisms of visual processing in monkeys. In **Handbook of neuropsychology, Volume 2** (F. Boller and J. Grafman, Eds). Amsterdam: Elsevier, pp. 267–299.

Dubrawski, A. and Crowley, J.L. (1994). Learning locomotion reflexes: A self-supervised neural system for a mobile robot. *Robotics and Autonomous Systems*, **12**, 133–142.

Ebner, T.J. and Bloedel, J.R. (1981). Correlation between activity of Purkinje cells and its modification by natural peripheral stimuli. *Journal of Neurophysiology*, **45**, 948–961.

Eichenbaum, H., Otto, T., and Cohen, N.J. (1994). Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences*, **17**, 449–472.

Estes, W.K. (1994). **Classification and cognition**. New York: Oxford University Press.

Fiala, J.C., Grossberg, S. and Bullock, D. (1995). Metabotropic glutamate receptor activation in cerebellar Purkinje cells as substrate for adaptive timing of the classically conditioned eye blink response. **Technical Report CAS/CNS-TR-95-029**. Boston, MA: Boston University.

Fujita, M. (1982a). Adaptive filter model of the cerebellum. *Biological Cybernetics*, **45**, 195–206.

Fujita, M. (1982b). Simulation of adaptive modification of the vestibulo-ocular reflex with an adaptive filter model of the cerebellum. *Biological Cybernetics*, **45**, 207–214.

Gaffan, D. (1974) *Journal of Comparative and Physiological Psychology*, **86**, 1100–1109.

Gaffan, D. (1985). Hippocampus: Memory, habit, and voluntary movement. *Philosophical Transactions of the Royal Society of London*, **B308**, 87–99.

Gaffan, D. (1992) *European Journal of Neuroscience*, **4**, 381–388.

Gaffan, D. (1994). Interaction of temporal lobe and frontal lobe in memory. In A.M. Thierry, J. Glowsinski, P.S. Goldman-Rakic, and Y. Christen (Eds.), **Motor and cognitive functions of the prefrontal cortex**. New York: Springer-Verlag, pp. 129–138.

Gilbert, P.F.C. and Thach, W.T. (1977). Purkinje cell activity during motor learning. *Brain Research*, **128**, 309–328.

Gjerdingen, R.O. (1990). Categorization of musical patterns by self-organizing neuron like networks. *Music Perception*, **7**, 339–370.

Gochin, P.M., Miller, E.K., Gross, C.G., and Gerstein, G.L. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Experimental Brain Research*, **84**, 505–516.

Goldman-Rakic, P.S. (1994). The issue of memory in the study of prefrontal function. In A.M. Thierry, J. Glowsinski, P.S. Goldman-Rakic, and Y. Christen (Eds.), **Motor and cognitive functions of the prefrontal cortex**. New York: Springer-Verlag, pp. 112-121.

Goodman, P., Kaburlasos, V., Egbert, D., Carpenter, G.A., Grossberg, S., Reynolds, J.H., Hammermeister, K., Marshall, G., and Grover, F. (1992). Fuzzy ARTMAP neural network prediction of heart surgery mortality. In G.A. Carpenter and S. Grossberg (Eds.), **Neural networks for learning, recognition, and control**. Tyngsboro, MA: Wang Institute of Boston University, p. 48.

Gormezano, I. and Kehoe, E.J. (1984). Associative transfer in classical conditioning to serial compounds. In: **Quantitative analyses of behavior, Volume 3: Acquisition**. Cambridge, MA: Ballinger, pp. 297–322,

Gove, A., Grossberg, S., and Mingolla, E. (1995). Brightness perception, illusory contours, and corticogeniculate feedback. **Technical Report CAS/CNS-TR-94-033**, Boston, MA: Boston University. *Visual Neuroscience*, in press.

Govindarajan, K.K., Grossberg, S., Wyse, L.L., and Cohen, M.A. (1994). A neural network model of auditory scene analysis and source segregation. **Technical Report CAS/CNS-TR-94-039**, Boston, MA: Boston University. Submitted for publication.

Graf, P., Squire, L.R., and Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **10**, 164–178.

Gray, J.A. (1982). **The neuropsychology of anxiety: An enquiry into the functions of the septo-hippocampal system**. New York: Oxford University Press.

Grossberg, S. (1968a). Some nonlinear networks capable of learning a spatial pattern of arbitrary complexity. *Proceedings of the National Academy of Sciences*, **59**, 368–372.

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

Grossberg, S. (1969a). On learning and energy-entropy dependence in recurrent and nonrecurrent signed networks. *Journal of Statistical Physics*, **1**, 319–350.

Grossberg, S. (1969b). On the production and release of chemical transmitters and related topics in cellular control. *Journal of Theoretical Biology*, **22**, 325–364.

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

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

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

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

Grossberg, S. (1972c). Neural expectation: Cerebellar and retinal analogs of cells fired by learnable or unlearned pattern classes. *Kybernetik*, **10**, 49–57.

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

Grossberg, S. (1974). Classical and instrumental learning by neural networks. In R. Rosen and F. Snell (Eds.), **Progress in theoretical biology**. New York: Academic Press. Reprinted in S. Grossberg, **Studies of mind and brain**. Boston: Reidel Press, 1982.

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

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

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

Grossberg, S. (1978). A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and plans. In R. Rosen and F. Snell (Eds.), **Progress in theoretical biology, Volume 5**. New York: Academic Press. Reprinted in S. Grossberg, **Studies of mind and brain**. Boston: Reidel Press, 1982.

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

Grossberg, S. (1982a). Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychological Review*, **89**, 529–572.

Grossberg, S. (1982b). **Studies of mind and brain**. Boston: Reidel Press.

Grossberg, S. (1984). Some psychological correlates of a developmental, cognitive and motivational theory. *Annals of the New York Academy of Sciences*, **425**, 58–142.

Grossberg, S. (1986). The adaptive self-organization of serial order in behavior: Speech, language, and motor control. In E.C. Schwab and H.C. Nusbaum (Eds.), **Pattern recognition by humans and machines, Vol. 1: Speech perception**. New York: Academic Press, pp.187–294.

Grossberg, S. (Ed.) (1987). **The adaptive brain, Volumes I and II**. Amsterdam: Elsevier/North-Holland.

Grossberg, S. (Ed.) (1988a). **Neural networks and natural intelligence**. Cambridge, MA: MIT Press.

Grossberg, S. (1988b). Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, **1**, 17–61.

Grossberg, S., Boardman, I., and Cohen, M.A. (1994). Neural dynamics of variable-rate speech categorization. **Technical Report CAS/CNS-TR-94-038**, Boston, MA: Boston University. Submitted for publication.

Grossberg, S. and Gutowski, W.E. (1987). Neural dynamics of decision making under risk: Affective balance and cognitive-emotional interactions. *Psychological Review*, **94**, 300–318.

Grossberg, S. and Kuperstein, M. (1986). **Neural dynamics of adaptive sensory-motor control: Ballistic eye movements**. Amsterdam: Elsevier/North-Holland.

Grossberg, S. and Kuperstein, M. (1989). **Neural dynamics of adaptive sensory-motor control: Expanded edition**. Elmsford, NY: Pergamon Press.

Grossberg, S. and Levine, D.S. (1987). Neural dynamics of attentionally modulated Pavlovian conditioning: Blocking, inter-stimulus interval, and secondary reinforcement. *Applied Optics*, **26**, 5015–5030.

Grossberg, S. and Merrill, J.W.L. (1992). A neural network model of adaptively timed reinforcement learning and hippocampal dynamics. *Cognitive Brain Research*, **1**, 3–38.

Grossberg, S. and Merrill, J.W.L. (1995). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. **Technical Report CAS/CNS-TR-93-065**. Boston, MA: Boston University. *Journal of Cognitive Neuroscience*, in press

Grossberg, S. and Pepe, J. (1970). Schizophrenia: Possible dependence of associational span, bowing, and primacy versus recency on spiking threshold. *Behavioral Science*, **15**, 359–362.

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

Grossberg, S. and Schmajuk, N.A. (1987). Neural dynamics of Pavlovian conditioning: Conditioned reinforcement, inhibition, and opponent processing. *Psychobiology*, **15**, 195–240.

Grossberg, S. and Schmajuk, N.A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, **2**, 79–102.

Grossberg, S. and Somers, D. (1991). Synchronized oscillations during cooperative feature linking in a cortical model of visual perception. *Neural Networks*, **4**, 453–466.

Grossberg, S. and Stone, G.O. (1986a). Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review*, **93**, 46–74.

Grossberg, S. and Stone, G.O. (1986b). Neural dynamics of attention switching and temporal order information in short term memory. *Memory and Cognition*, **14**, 451–468.

Ham, F. and Han, S. (1993). Quantitative study of ARS complex using fuzzy ARTMAP and MIT/BIH arrythmia database. In **Proceedings of the world congress on neural networks, I**, 207–211. Hillsdale, NJ: Erlbaum Associates.

Harries, M.H. and Perrett, D.I. (1991). Visual processing of faces in temporal cortex: Physiological evidence for a modular organization and possible anatomical correlates. *Journal of Cognitive Neuroscience*, **3**, 9–24.

Harvey, R.M. (1993). Nursing diagnostics by computers: An application of neural networks. *Nursing Diagnostics*, **4**, 26–34.

Henneman, E. (1957). Relation between size of neurons and their susceptibility to discharge. *Science*, **26**, 1345–1347.

Henneman, E. (1985). The size-principle: A deterministic output emerges from a set of probabilistic connections. *Journal of Experimental Biology*, **115**, 105–112.

Hoehler, F.K. and Thompson, R.F. (1980). Effects of the interstimulus (CS–UCS) interval on hippocampal unit activity during classical conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*). *Journal of Comparative and Physiological Psychology*, **94**, 201–215.

Holder, M.D. and Roberts, S. (1985). Comparison of timing and classical conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **11**, 172–193.

Ito, M. (1984). **The cerebellum and neural control**. New York: Raven Press.

Jones, R. and Keck, M.J. (1978). Visual evoked response as a function of grating spatial frequency. *Investigative Ophthalmology and Visual Science*, **17**, 652–659.

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

Kandel, E.R. and O'Dell, T.J. (1992). Are adult learning mechanisms also used for development? *Science*, **258**, 243–245.

Kasperkiewicz, J., Racz, J., and Dubrawski, A. (1994). HPC strength prediction using artificial neural networks for development of diagnostic monitoring system in nuclear plants. *ASCE Journal of Computing in Civil Engineering*.

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

Kehoe, E.J., Gibbs, C.M., Garcia, E., and Gormenzano, I. (1979). Associative transfer and stimulus selection in classical conditioning of the rabbit's nictitating membrane response to

serial compound CS's. *Journal of Experimental Psychology: Animal Behavior Processes*, **5**, 1–57.

Kehoe, E.J., Horne, P.S., Macrae, M., and Horne, A.J. (1993) *Journal of Experimental Psychology: Animal Behavior Processes*, **19**, 265–283.

Kehoe, E.J., Marshall-Goodell, B., and Gormenzano, I. (1987). Differential conditioning of the rabbit's nictitating membrane response to serial compound stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **13**, 17–30.

Kehoe, E.J. and Morrow, L.D. (1984). Temporal dynamics of the rabbit's nictitating membrane response in serial compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 205–220.

Keyvan, S., Durg, A., and Rabelo, L. (1993). Application of artificial neural networks for development of diagnostic monitoring system in nuclear plants. *Transactions of the American Nuclear Society*, **1**, 515–522.

Knight, R.T. (1994). Attention regulation and human prefrontal cortex. In A.M. Thierry, J. Glowsinksi, P.S. Goldman-Rakic, and Y. Christen (Eds.), **Motor and cognitive functions of the prefrontal cortex**. New York: Springer-Verlag, pp. 160–173.

Knowlton, B.J. and Squire, L.R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, **262**, 1747–1749.

Kohonen, T. (1984). **Self-organization and associative memory**. New York: Springer-Verlag.

Kuno, M. (1995). **The synapse: Function, plasticity, and neurotrophism**. Oxford: Oxford University Press.

Leonard, B. and McNaughton, B.L. (1990). Spatial representation in the rat: Conceptual, behavioral, and neurophysiological perspectives. In **Neurobiology of comparative cognition** (R.P. Kesner and D.S. Olton, Eds.). Hillsdale, NJ: Erlbaum Associates, pp. 363–422.

Levy, W.B. (1985). Associative changes at the synapse: LTP in the hippocampus. In **Synaptic modification, neuron selectivity, and nervous system organization** (W.B. Levy, J. Anderson, and S. Lehmkuhle, Eds.). Hillsdale, NJ: Erlbaum Associates, pp. 5–33.

Levy, W.B. and Desmond, N.L. (1985). The rules of elemental synaptic plasticity. In **Synaptic modification, neuron selectivity, and nervous system organization** (W.B. Levy, J. Anderson, and S. Lehmkuhle, Eds.). Hillsdale, NJ: Erlbaum Associates, pp. 105–121.

Lisberger, S.G. (1988). The neural basis for motor learning in the vestibulo-ocular reflex in monkeys. *Trends in Neurosciences*, **11**, 147–152.

Logothetis, N., Pauls, J., Buelthoff, H., and Poggio, T. (1994). View-dependent object recognition by monkeys. *Current Biology*, **4**, 401–414.

Lynch, G., McGaugh, J.L., and Weinberger, N.M. (Eds.) (1984). **Neurobiology of learning and memory**. New York: Guilford Press.

Malsburg, C. von der (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, **14**, 85–100.

Maricq, A.V., Roberts, S., and Church, R.M. (1981). Methamphetamine and time estimation. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 18–30.

Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology*, **202**, 437–470.

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London*, **B-262**, 23–81.

McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1994). Why are there complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. **Technical Report PDP.CNS.94.1**, Pittsburgh, Pa: Carnegie Mellon University.

Meck, W.H. and Church, R.M. (1987). Cholinergic modulation of the content of temporal memory. *Behavioral Neuroscience*, **101**, 457–464.

Metha, B., Vij, L., and Rabelo, L. (1993). Prediction of secondary structures of proteins using fuzzy ARTMAP. In **Proceedings of the world congress on neural networks, I**, 228–232. Hillsdale, NJ: Erlbaum Associates.

Millenson, J.R., Kehoe, E.J., and Gormenzano, I. (1977). Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS-US intervals. *Learning and Motivation*, **8**, 351–366.

Miller, E.K., Li, L., and Desimone, R. (1991). A neural mechanism for working and recognition memory in inferior temporal cortex. *Science*, **254**, 1377–1379.

Miller, G.A. (1956). The magic number seven plus or minus two. *Psychological Review*, **63**, 81–97.

Milner, P. (1989). A cell assembly theory of hippocampal amnesia. *Neuropsychologia*, **27**, 23–30.

Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not separate removal of the amygdala and hippocampus. *Nature*, **273**, 297–298.

Mishkin, M. (1982). A memory system in the monkey. *Philosophical Transactions of the Royal Society of London B*, **298**, 85–95.

Mishkin, M. (1993). Cerebral memory circuits. In **Exploring brain functions: Models in neuroscience** (T.A. Poggio and D.A. Glaser, Eds.). New York: Wiley and Sons, pp. 113–125.

Mishkin, M. and Appenzeller, T. (1987). The anatomy of memory. *Scientific American*, **256**, 80–89.

Mishkin, M. and Delacour, J. (1975) *Journal of Experimental Psychology: Animal Behavior Processes*, **1**, 326–334.

Mishkin, M., Malamut, B., and Bachevalier, J. (1984). Memories and habits: Two neural systems. In **The neurobiology of learning and memory** (J.L. McGaugh, G. Lynch, and N. Weinberger, Eds.). New York: Guilford Press, pp. 287–296.

Moya, M.M., Koch, M.W., and Hostetler, L.D. (1993). One-class classifier networks for target recognition applications. In **Proceedings of the world congress on neural networks, III**, 797–801. Hillsdale, NJ: Erlbaum Associates.

Musselwhite, M.J. and Jeffreys, D.A. (1985). The influence of spatial frequency on the reaction times and evoked potentials recorded to grating pattern stimuli. *Vision Research*, **25**, 1545–1555.

Nowak, A.J. and Berger, T.W. (1992). Functional three-dimensional distribution of entorhinal projections to dentate granule cells of the *in vivo* rabbit hippocampus. *Society for Neuroscience Abstracts*, **18**, p. 321, #141.15.

O'Keefe, J. (1990). A computational theory of the hippocampal cognitive map. In **Progress in brain research** (J. Storm-Mathisen, J. Zimmer, and O.P. Ottersen, Eds). Amsterdam: Elsevier, pp. 301–312.

O'Keefe, J. and Nadel, L. (1978). **The hippocampus as a cognitive map**. New York: Oxford University Press.

Optican, L.M. and Robinson, D.A. (1980). Cerebellar-dependent adaptive control of primate saccadic system. *Journal of Neurophysiology*, **44**, 108–176.

Otto, T. and Eichenbaum, H. (1992). Neuronal activity in the hippocampus during delayed non-match to sample performance in rats: Evidence for hippocampal processing in recognition memory. *Hippocampus*, **2**, 323–334.

Parker, D.M. and Salzen, E.A. (1977a). Latency changes in the human visual evoked response to sinusoidal gratings. *Vision Research*, **17**, 1201–1204.

Parker, D.M. and Salzen, E.A. (1977b). The spatial selectivity of early and late waves within the human visual evoked response. *Perception*, **6**, 85–95.

Parker, D.M., Salzen, E.A., and Lishman, J.R. (1982a). Visual-evoked responses elicited by the onset and offset of sinusoidal gratings: Latency, waveform, and topographic characteristics. *Investigative Ophthalmology and Visual Sciences*, **22**, 657–680.

Parker, D.M., Salzen, E.A., and Lishman, J.R. (1982b). The early waves of the visual evoked potential to sinusoidal gratings: Responses to quadrant stimulation as a function of spatial frequency. *Electroencephalography and Clinical Neurophysiology*, **53**, 427–435.

Perrett, D.I., Mistlin, A.J., and Chitty, A.J. (1987). Visual cells responsive to faces. *Trends in Neurosciences*, **10**, 358–364.

Perrett, S.P., Ruiz, B.P., and Mauk, M.D. (1993). Cerebellar cortex lesions disrupt learning-dependent timing of conditioned eyelid responses. *Journal of Neuroscience*, **13**, 1708–1718.

Plant, G.T., Zimmern, R.L., and Durden, K. (1983). Transient visually evoked potentials to the pattern reversal and onset of sinusoidal gratings. *Electroencephalography and Clinical Neurophysiology*, **56**, 147–158.

Posner, M.I. and Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, **77**, 353–363.

Posner, M.I. and Keele, S.W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, **83**, 304–308.

Pribram, K.H. (1986). The hippocampal system and recombinant processing. In **The hippocampus, Volume 4** (R.L. Isaacson and K.H. Pribram, Eds.). New York: Plenum Press, pp. 329–370.

Rauschecker, J.P. and Singer, W. (1979). Changes in the circuitry of the kitten's visual cortex are gated by postsynaptic activity. *Nature*, **280**, 58–60.

Rickert, E.J., Bennett, T.L., Lane, P.L., and French, J. (1978). Hippocampectomy and the attenuation of blocking. *Behavioral Biology*, **22**, 147–160.

Roberts, K., Aguilar, J.M., Bullock, D., and Grossberg, S. (1994). Superior colliculus dynamics reflect interactions between planned and reactive saccades. *Investigative Ophthalmology and Visual Science*, **35**, #3613.

Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 242–268.

Ryle, G. (1949). **The concept of mind**. Hutchinson Press.

Schmajuk, N.A., Spear, N.E., and Isaacson, R.L. (1983). Absence of overshadowing in rats with hippocampal lesions. *Physiological Psychology*, **11**, 59–62.

Schwartz, E.L., Desimone, R., Albright, T., and Gross, C.G. (1983). Shape recognition and inferior temporal neurons. *Proceedings of the National Academy of Sciences*, **80**, 5776–5778.

Scoville, W.B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesion. *Journal of Neurology, Neurosurgery, and Psychiatry*, **20**, 11–21.

Seibert, M. and Waxman, A.M. (1991). Learning and recognizing 3-D objects from multiple views in a neural system. In H. Wechsler (Ed.), **Neural networks for perception, Volume 1**. New York: Academic Press.

Seibert, M. and Waxman, A.M. (1992). Adaptive 3-D object recognition from multiple views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**, 107–124.

Sillito, A.M., Jones, H.E., Gerstein, G.L., and West, D.C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, **369**, 479–482.

Singer, W. (1983). Neuronal activity as a shaping factor in the self-organization of neuron assemblies. In **Synergetics of the brain** (E. Basar, H. Flohr, H. Haken, and A.J. Mandell,

Eds.). New York: Springer-Verlag, pp. 89–101.

Skrandies, W. (1984). Scalp potential fields evoked by grating stimuli: Effects of spatial frequency and orientation. *Electroencephalography and Clinical Neurophysiology*, **58**, 325–332.

Smith, M.C. (1968). CS–US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, **3**, 679–687.

Solomon, P.R. (1977). The role of hippocampus in blocking and conditioned inhibition of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, **91**, 407–417.

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, **240**, 338–340.

Squire, L.R. and Butters, N. (Eds.) (1984). **Neuropsychology of memory**. New York: Guilford Press.

Squire, L.R. and Cohen, N.J. (1984). Human memory and amnesia. In **Neurobiology of learning and memory** (G. Lynch, J. McGaugh, and N.M. Weinberger, Eds.). New York: Guilford Press, pp. 3–64.

Squire, L.R. and Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, **253**, 1380–1386.

Staddon, J.E.R. (1983). **Adaptive behavior and learning**. New York: Cambridge University Press.

Strouthes, A. (1965). Effect 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*, **69**, 287–291.

Suzuki, Y. (1995). Self-organizing QRS-wave recognition in ECG using neural networks. *IEEE Transactions on Neural Networks*, in press.

Suzuki, Y., Abe, Y., and Ono, K. (1994). Self-organizing QRS wave recognition system in ECG using ART 2. In **Proceedings of the world congress on neural networks, IV**, 39–42. Hillsdale, NJ: Erlbaum Associates.

Thompson, R.F. (1988). The neural basis of basic associative learning of discrete behavioral responses. *Trends in Neurosciences*, **11**, 152–155.

Thompson, R.F., Barchas, J.D., Clark, G.A., Donegan, N., Kettner, R.E., Lavond, D.G., Madden, J., Mauk, M.D., and McCormick, D.A. (1984). Neuronal substrates of associative learning in the mammalian brain. In **Primary neural substrates of learning and behavioral change** (D.L. Aldon and J. Farley, Eds.). New York: Cambridge University Press, pp. 71–99.

Thompson, R.F., Clark, G.A., Donegan, N.H., Lavond, G.A., Lincoln, D.G., Maddon, J., Mamounas, L.A., Mauk, M.D., and McCormick, D.A. (1987). Neuronal substrates of discrete, defensive conditioned reflexes, conditioned fear states, and their interactions in the rabbit. In **Classical conditioning, Third edition** (I. Gormenzano, W.F. Prokasy, and R.F. Thompson, Eds). Hillsdale, NJ: Erlbaum Associates, pp. 371–399.

Vassilev, A., Manahilov, V., and Mitov, D. (1983). Spatial frequency and pattern onset-offset response. *Vision Research*, **23**, 1417–1422.

Vassilev, A. and Strashimirov, D. (1979). On the latency of human visually evoked response to sinusoidal gratings. *Vision Research*, **19**, 843–846.

Warrington, E.K. and Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychology*, **12**, 419–428.

Wienke, D., Xie, P., and Hopke, P.K. (1994). An adaptive resonance theory based artificial neural network (ART 2A) for rapid identification of airborne particle shapes from their

scanning electron microscopy images. *Chemometrics and Intelligent Laboratory Systems.*

Wilkie, D.M. (1987). Stimulus intensity affects pigeons' timing behavior: Implications for an internal clock model. *Animal Learning and Behavior*, **15**, 35–39.

Williamson, S.J., Kaufman, I., and Brenner, D. (1978). Latency of the neuromagnetic response of the human visual cortex. *Vision Research*, **18**, 107–110.

Zola-Morgan, S.M. and Squire, L.R. (1990). The primate hippocampal formation: Evidence for a time-limited role in memory storage. *Science*, **250**, 288–290.

FIGURE CAPTIONS

Figure 1. An example of a model ART circuit in which attentional and orienting circuits interact. Level $\mathcal{F}_1$ encodes a distributed representation of an event by a short term memory (STM) activation pattern across a network of feature detectors. Level $\mathcal{F}_2$ encodes the event using a compressed STM representation of the $\mathcal{F}_1$ pattern. Learning of these recognition codes occurs at the long term memory (LTM) traces within the bottom-up and top-down pathways between levels $\mathcal{F}_1$ and $\mathcal{F}_2$. The top-down pathways read-out learned expectations whose prototypes are matched against bottom-up input patterns at $\mathcal{F}_1$. The size of mismatches in response to novel events are evaluated relative to the vigilance parameter ρ of the orienting subsystem $\mathcal{A}$. A large enough mismatch resets the recognition code that is active in STM at $\mathcal{F}_2$ and initiates a memory search for a more appropriate recognition code. Output from subsystem $\mathcal{A}$ can also trigger an orienting response.

Figure 2. ART search for a recognition code: (A) The input pattern **I** is instated across the feature detectors at level $\mathcal{F}_1$ as a short term memory (STM) activity pattern **X**. Input **I** also nonspecifically activates the orienting subsystem $\mathcal{A}$; see Figure 1. STM pattern **X** is represented by the hatched pattern across $\mathcal{F}_1$. Pattern **X** both inhibits $\mathcal{A}$ and generates the output pattern **S**. Pattern **S** is multiplied by long term memory (LTM) traces and added at $\mathcal{F}_2$ nodes to form the input pattern **T**, which activates the STM pattern **Y** across the recognition categories coded at level $\mathcal{F}_2$. (B) Pattern **Y** generates the top-down output pattern **U** which is multiplied by top-down LTM traces and added at $\mathcal{F}_1$ nodes to form the prototype pattern **V** that encodes the learned expectation of the active $\mathcal{F}_2$ nodes. If **V** mismatches **I** at $\mathcal{F}_1$, then a new STM activity pattern $\mathbf{X}^*$ is generated at $\mathcal{F}_1$. $\mathbf{X}^*$ is represented by the hatched pattern. It includes the features of **I** that are confirmed by **V**. Inactivated nodes corresponding to unconfirmed features of **X** are unhatched. The reduction in total STM activity which occurs when **X** is transformed into $\mathbf{X}^*$ causes a decrease in the total inhibition from $\mathcal{F}_1$ to $\mathcal{A}$. (C) If inhibition decreases sufficiently, $\mathcal{A}$ releases a nonspecific arousal wave to $\mathcal{F}_2$, which resets the STM pattern **Y** at $\mathcal{F}_2$. (D) After **Y** is inhibited, its top-down prototype signal is eliminated, and **X** can be reinstated at $\mathcal{F}_1$. Enduring traces of the prior reset lead **X** to activate a different STM pattern $\mathbf{Y}^*$ at $\mathcal{F}_2$. If the top-down prototype due to $\mathbf{Y}^*$ also mismatches **I** at $\mathcal{F}_1$, then the search for an appropriate $\mathcal{F}_2$ code continues until a more appropriate $\mathcal{F}_2$ representation is selected. Then an attentive resonance develops and learning of the attended data is initiated. [Reprinted with permission from Carpenter and Grossberg (1993).]

Figure 3. Schematic conditioning circuit: Conditioned stimuli (CS_i) activate sensory categories (S_{CS_i}) which compete among themselves for limited capacity short-term memory activation and storage, as at level $\mathcal{F}_2$ in an ART circuit. The activated S_{CS_i} representations elicit trainable signals to drive representations D and motor command representations M, denoted "response learning". Learning from a sensory representation S_{CS_i} to a drive representation D is called conditioned reinforcer learning. Learning from D to a S_{CS_i} is called incentive motivational learning. Signals from D to S_{CS_i} are elicited when the combination of conditioned sensory plus internal drive inputs is sufficiently large. Sensory representations that win the competition in response to the balance of external inputs and internal motivational signals can activate motor command pathways.

Figure 4. A spectrally timed conditioning model with feedback pathways $D \rightarrow S^{(2)} \rightarrow S^{(1)}$ that are capable of focusing attention in an adaptively timed fashion on reinforcing events. The sensory representations S of Figure 3 are here broken into two successive levels $S^{(1)}$ and $S^{(2)}$. Levels $S^{(1)}$ and $S^{(2)}$ interact via reciprocal excitatory pathways. The excitatory pathways $S^{(1)} \rightarrow D$ and $D \rightarrow S^{(2)}$ are, as in Figure 3, adaptive. Representations in $S^{(2)}$ can, however, fire only if they receive convergent signals from $S^{(1)}$ and D. Then they deliver positive feedback to $S^{(1)}$ and bias the competition to focus attention on their respective features and to attentionally block inhibited features. Prior to conditioning, a CS can only be stored in STM at $S^{(1)}$ and can subliminally prime $S^{(2)}$ and D representations without supraliminally firing these representations. After conditioning, the CS can trigger strong conditioned $S^{(1)} \rightarrow D \rightarrow S^{(2)} \rightarrow S^{(1)}$ feedback and rapidly draw attention to itself as it activates the emotional representations and motivational pathways controlled by D. Representation D can also inhibit the orienting subsystem as it focuses attention upon motivationally valued sensory events. The sensory representations $S^{(1)}$ send parallel pathways to a spectral timing circuit T whose adaptive weights z sample the Now Print, or teaching signal, N that is transiently activated by changes in the activity of the drive representation D. After conditioning of T takes place, adaptively timed readout from T can maintain attention on task-relevant cues for a learned duration via the $T \rightarrow D \rightarrow S$ feedback pathway. Timed signals also inhibit the orienting subsystem via the $T \rightarrow D \rightarrow \mathcal{A}$ pathway and thereby help to prevent distracting events from interfering with planned consummatory acts. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 5. A computer simulation of spectral timing: (a) In response to a CS input I_i in Figure 4, a spectrum of population activities x_{ij} react at different rates and generate signals $f_{ij} = f(x_{ij})$; (b) each signal causes a transmitter y_{ij} in its pathway to become inactivated, or habituate, at a different rate; (c) the transmitters y_{ij} multiply, or gate, the signals f_{ij} to generate net signals $g_{ij} = f_{ij}y_{ij}$ that sample overlapping time intervals; (d) the sampling signals g_{ij} and the US, expressed via the teaching signal N, conjointly activate adaptive weights, or LTM traces, z_{ij}, which generate adaptively gated output signals $h_{ij} = g_{ij}z_{ij}$; (e) although individual signals h_{ij} do not well time the ISI, the population sum $R = \sum_j h_{ij}$ of the adaptive signals does accurately time the ISI (dotted vertical lines). Parameters and signal functions: $\gamma = 0.2, \alpha_y = 1.0, \beta_y = 125.0, \alpha_z = 1.0, \delta = 0.0, \epsilon = 0.02, \alpha_E = 240.0, \alpha_A = 1.2, \beta_A = 120.0, \gamma_A = 12.0, \alpha_D = 120.0, \beta_D = 120.0, \gamma_D = 0.0, f_D(S) = [S - 0.05]^+, \alpha_C = 0.5, \beta_C = 25.0, f_C(D) = [D - 0.05]^+, f_A(A) = [A - 0.1]^+, F_X(A) = [A - 0.7]^+, r_j = 10.125 \ / \ (0.0125 + j)$; and the intensities of the CS and US inputs I_i in (1) equal 2. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 6. A model of adaptively timed cerebellar conditioning: US-activated climbing fibers provide a teaching signal that causes adaptively timed long term depression at parallel fiber-Purkinje cell synapses, thereby disinhibiting the inhibitory effect of tonic Purkinje cell outputs on cerebellar nuclear cells. The climbing fibers also control learning of adaptive gains along subcortical pathways through the nuclear cells. The net effect of learning is to open an adaptively timed Purkinje gate that enables learned gains to be expressed at the correct time.

Figure 7. In a START model, STM storage of a brief CS is achieved by positive feedback within the sensory representation S. CS attenuation by the US is dynamically controlled by the strength of recurrent inhibitory signals. (a) Input I_1 activated by CS_1; (b) Input I_0 activated by US; (c) STM activation of CS_1 sensory representation; (d) STM activation of US sensory representation. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 8. Behavior of the Now Print module of the START model after many conditioning trials: (a) Activation of the sensory representation S_1 by the CS; (b) Activation of the sensory representation S_0 by the US; (c) The resultant activation D of the drive representation D; (d) The resultant Now Print signal N. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 9. Generation of a Now Print signal: (a) The output of a drive representation D is converted into a Now Print signal N by passing this output through a fast excitatory pathway and a slower inhibitory pathway, whose signals converge at N; (b) Simulation of the activity D of D in response to two successive inputs, with the first response larger; (c) Activity N of scales with the size of the increment in D. All parameters were as in Figure 5. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 10. Selective learning within different spectral populations at a fixed ISI = 0.5 time units. Each three-image panel from (a) to (f) represents the gated signal $g_{1j}(t)$ [top], long term memory trace $z_{1j}(t)$ [middle], and doubly gated signal $h_{1j}(t) = g_{1j}(t)z_{1j}(t)$ [bottom], at a different value of j. In (a), $j = 1$; in (b), $j = 17$; in (c), $j = 33$; in (d), $j = 49$; in (e), $j = 65$; in (f), $j = 81$. The same parameters as in Figure 5 were used. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 11. The effect of self-printing upon the output of the model. (a) A large threshold ϵ in the Now Print signal abolishes self-printing and secondary conditioning. It generates the lower output $R(t)$. (b) A smaller threshold allows secondary conditioning and self-printing without a loss of timing. It generates the larger output $R(t)$. [Reprinted with permission from Grossberg and Merrill (1992).]

Figure 12. START model output $R(t)$ during secondary conditioning with varying ISIs between the first and second CS, and between the second CS and the US. Notation ISI_1 below denotes the ISI between CS_1 and US, and ISI_2 denotes the ISI between CS_2 and CS_1. On each learning trial either CS_1–US or CS_2–CS_1 occur, but not CS_2–CS_1–US. The curves are drawn with CS_1–US pairings in the left column and CS_2–CS_1 pairings in the right column. The vertical bars occur at successive 0.25 time unit intervals: (a),(b) $ISI_1 = .25$, $ISI_2 = .25$; (c),(d) $ISI_1 = .5$, $ISI_2 = .25$; (e),(f) $ISI_1 = .25$, $ISI_2 = .5$; (g),(h) $ISI_1 = .5$, $ISI_2 = .5$. [Reprinted with permission from Grossberg and Merrill (1992).]

Table 1

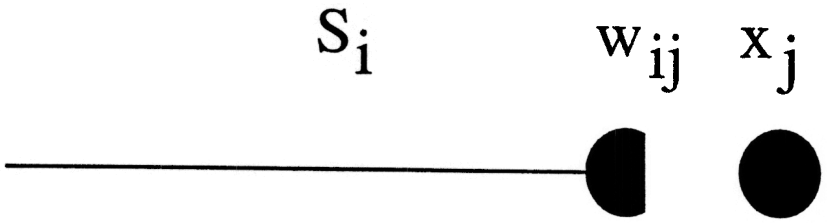

	Case 1	Case 2	Case 3	Case 4
State of S_i	+	−	+	−
State of x_j	+	+	−	−
State of w_{ij}	↑	↓	↔	↔

+ = active

− = inactive

↑ = increase

↓ = decrease

↔ = no change

Table 1: The instar learning, or gated steepest descent learning rule, embodies both Hebbian (LTP) and anti-Hebbian (LTD) properties within a single process.

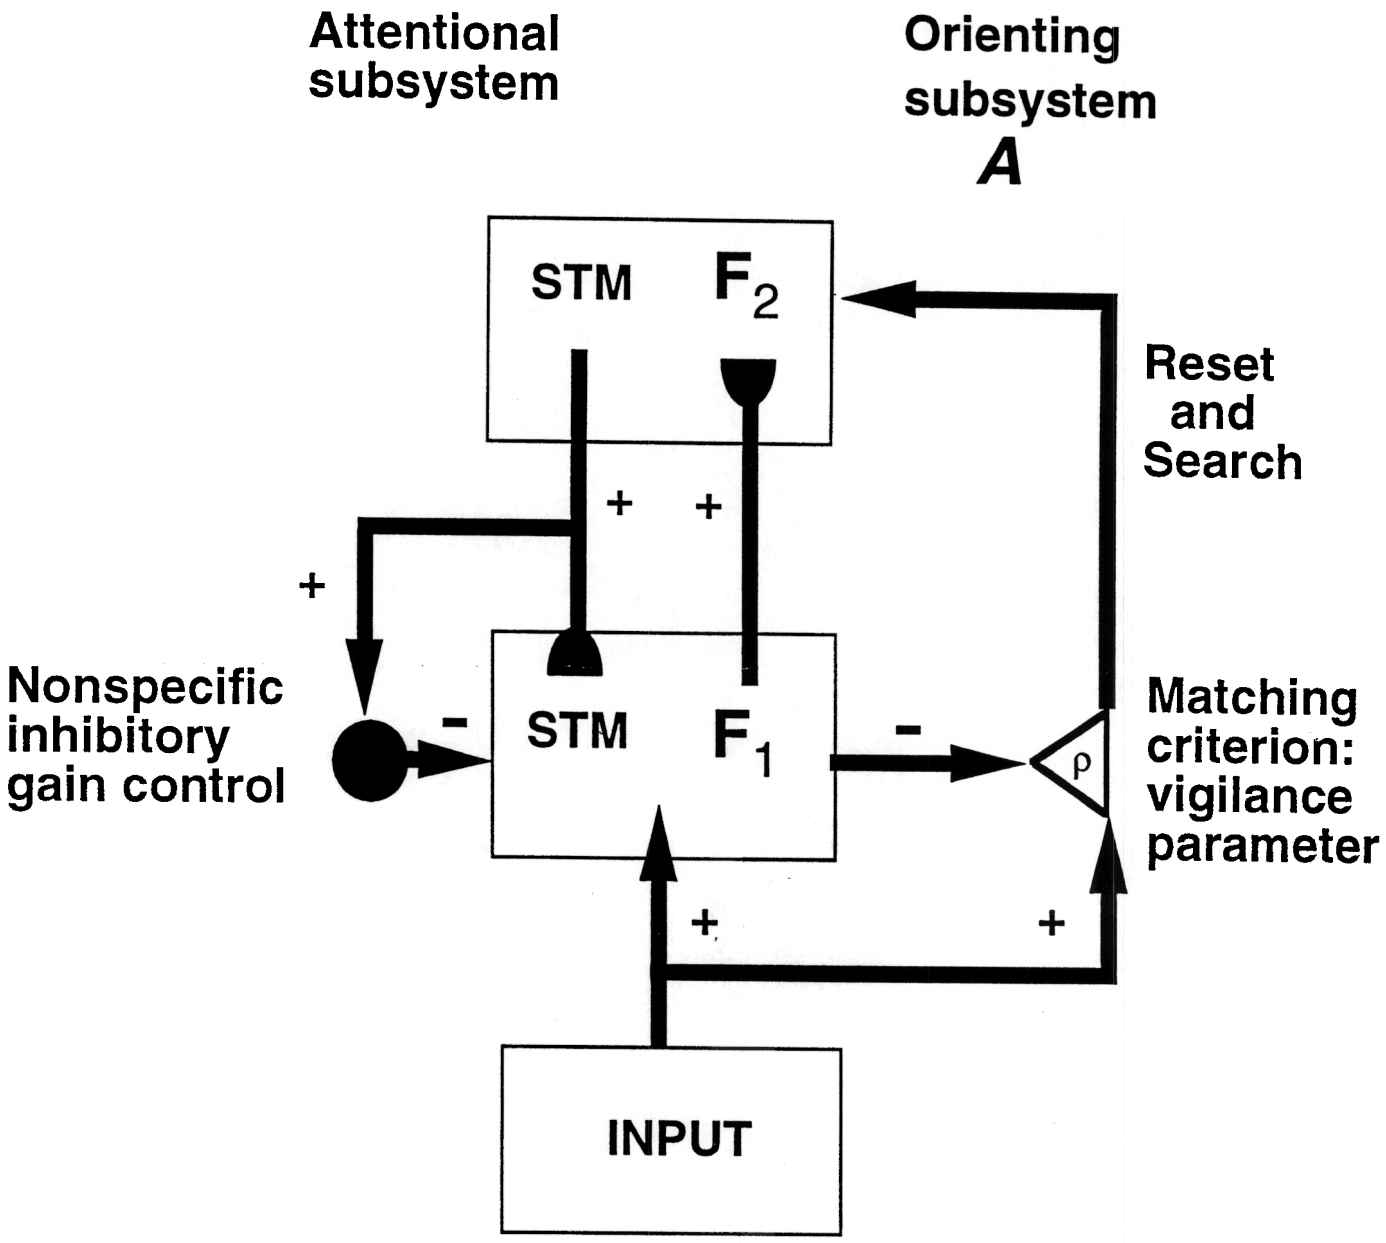
Attentional subsystem
Orienting subsystem
A
STM
F_2
Reset and Search
+
+
+
Nonspecific inhibitory gain control
-
STM
F_1
-
ρ
Matching criterion: vigilance parameter
+
+
INPUT

Figure

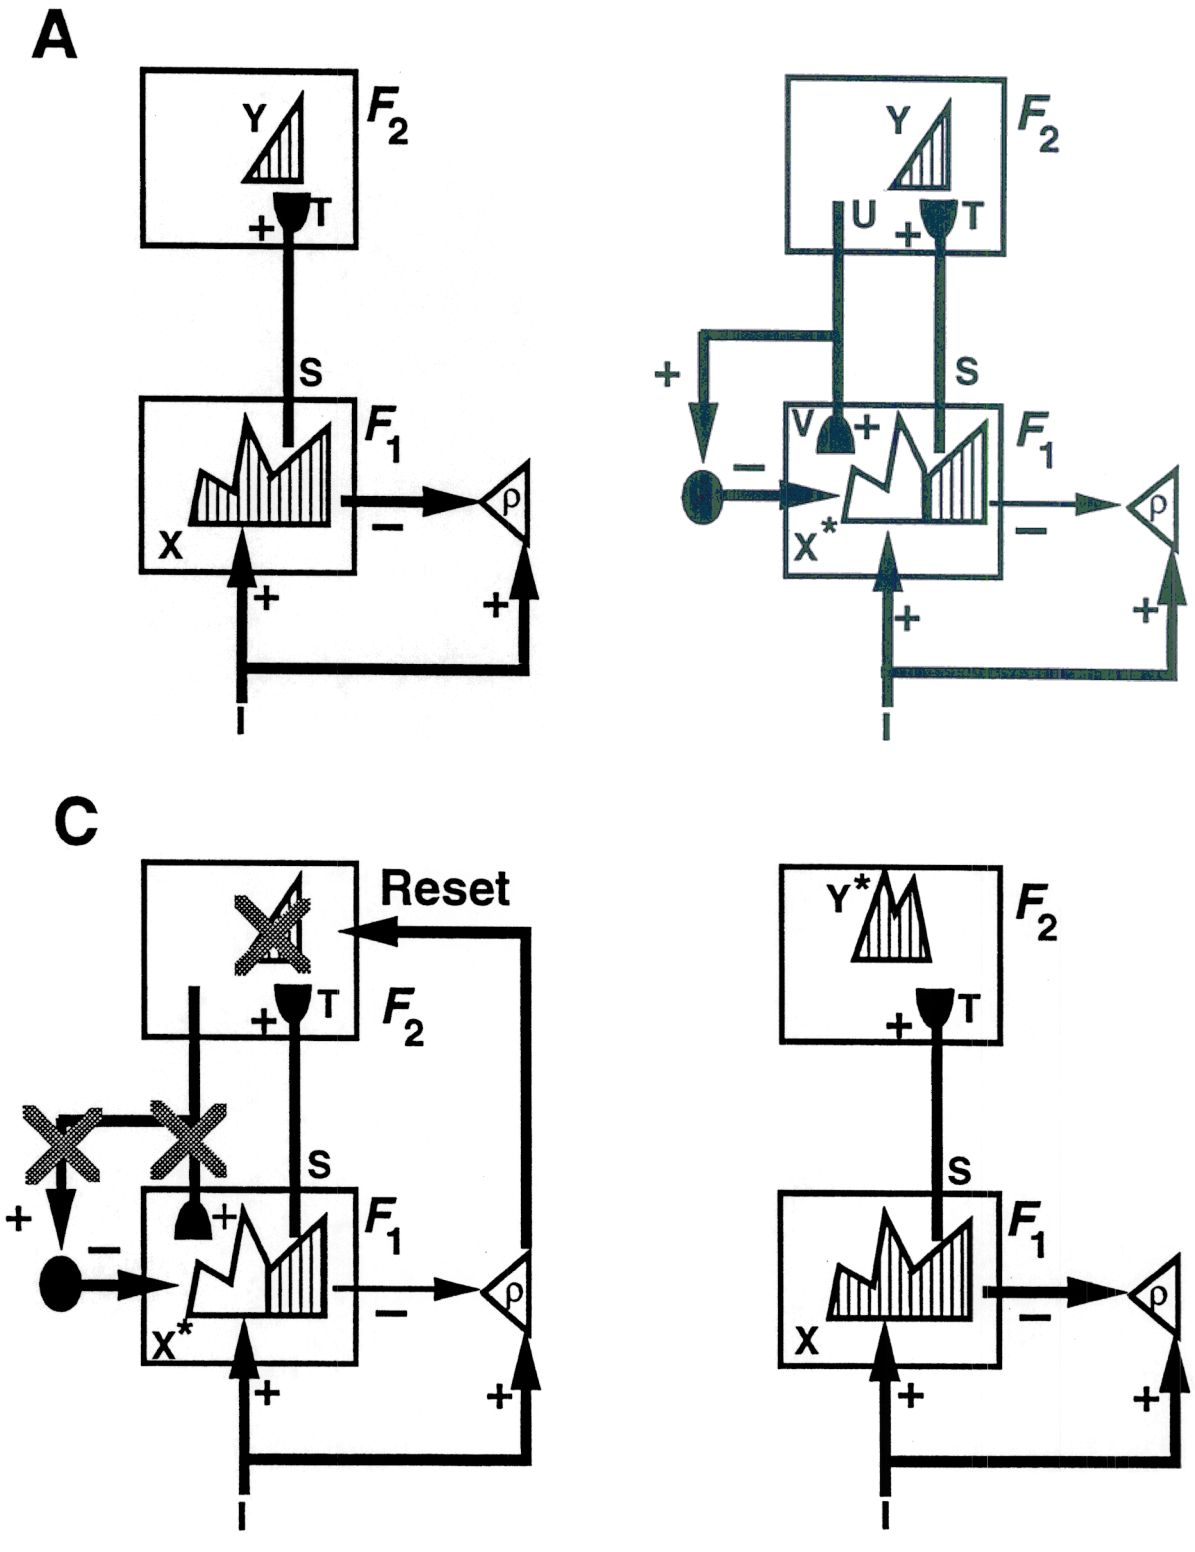

Figure

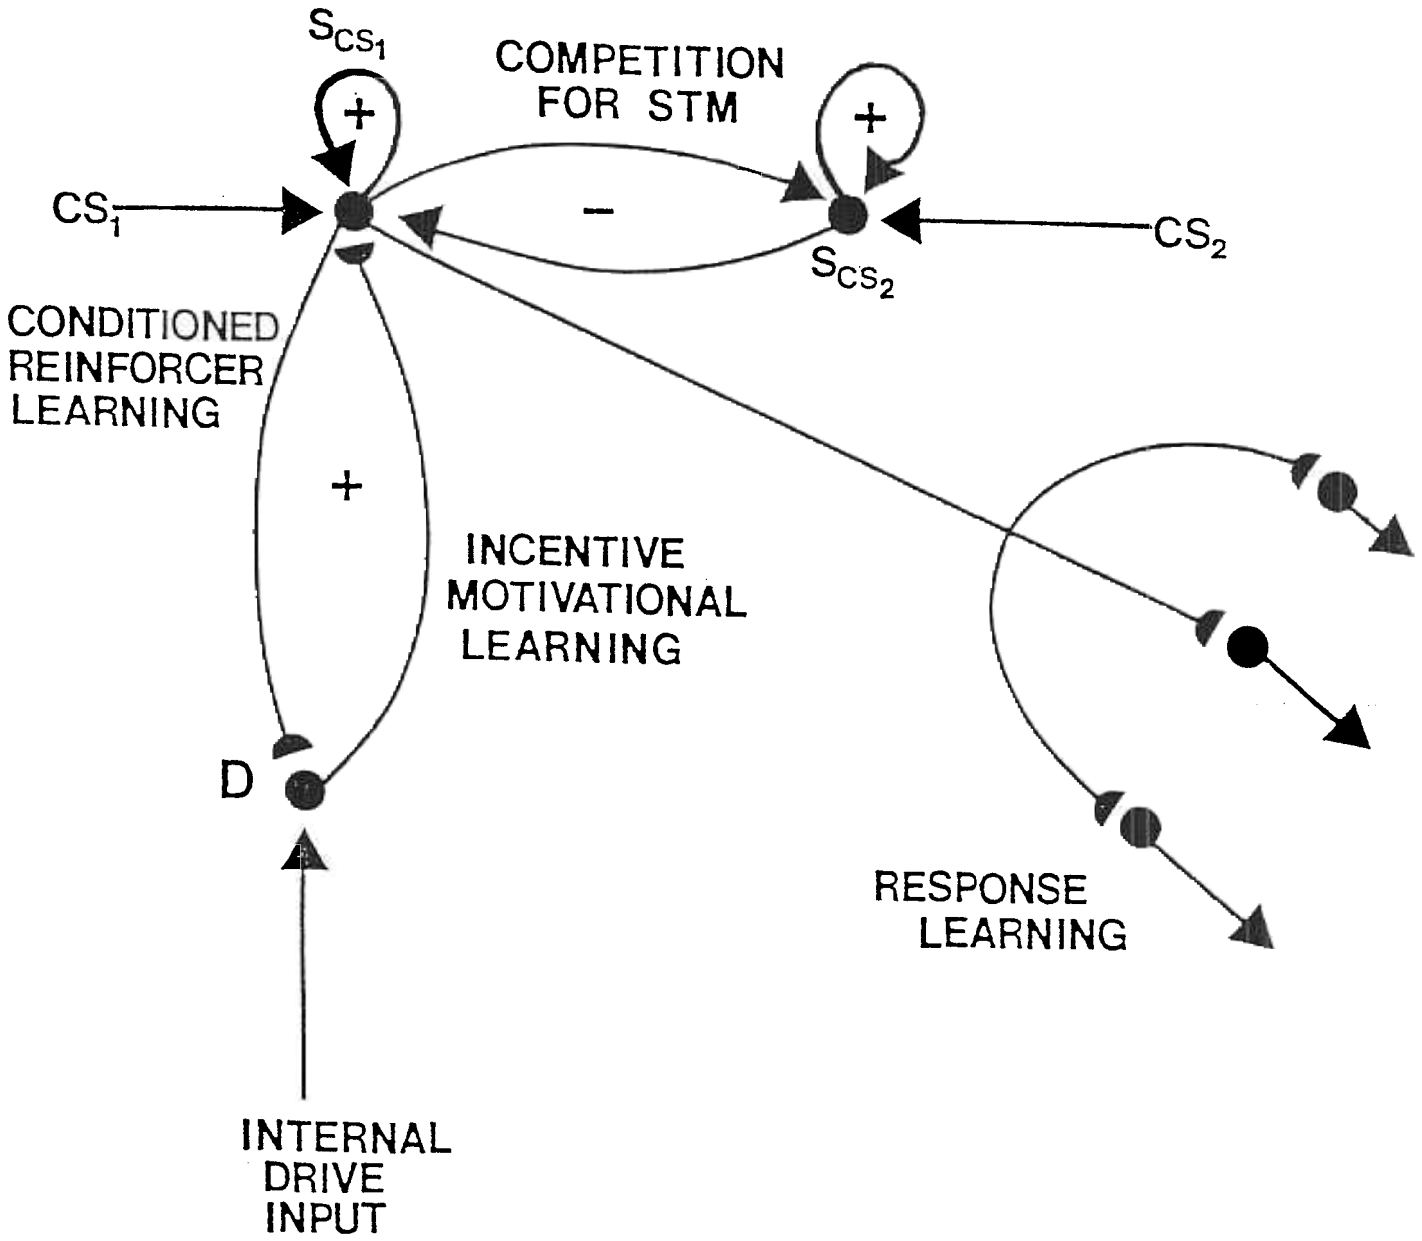
S_{CS_1}
COMPETITION FOR STM
+
+
CS_1
−
CS_2
S_{CS_2}
CONDITIONED REINFORCER LEARNING
+
INCENTIVE MOTIVATIONAL LEARNING
D
RESPONSE LEARNING
INTERNAL DRIVE INPUT

Figure 3

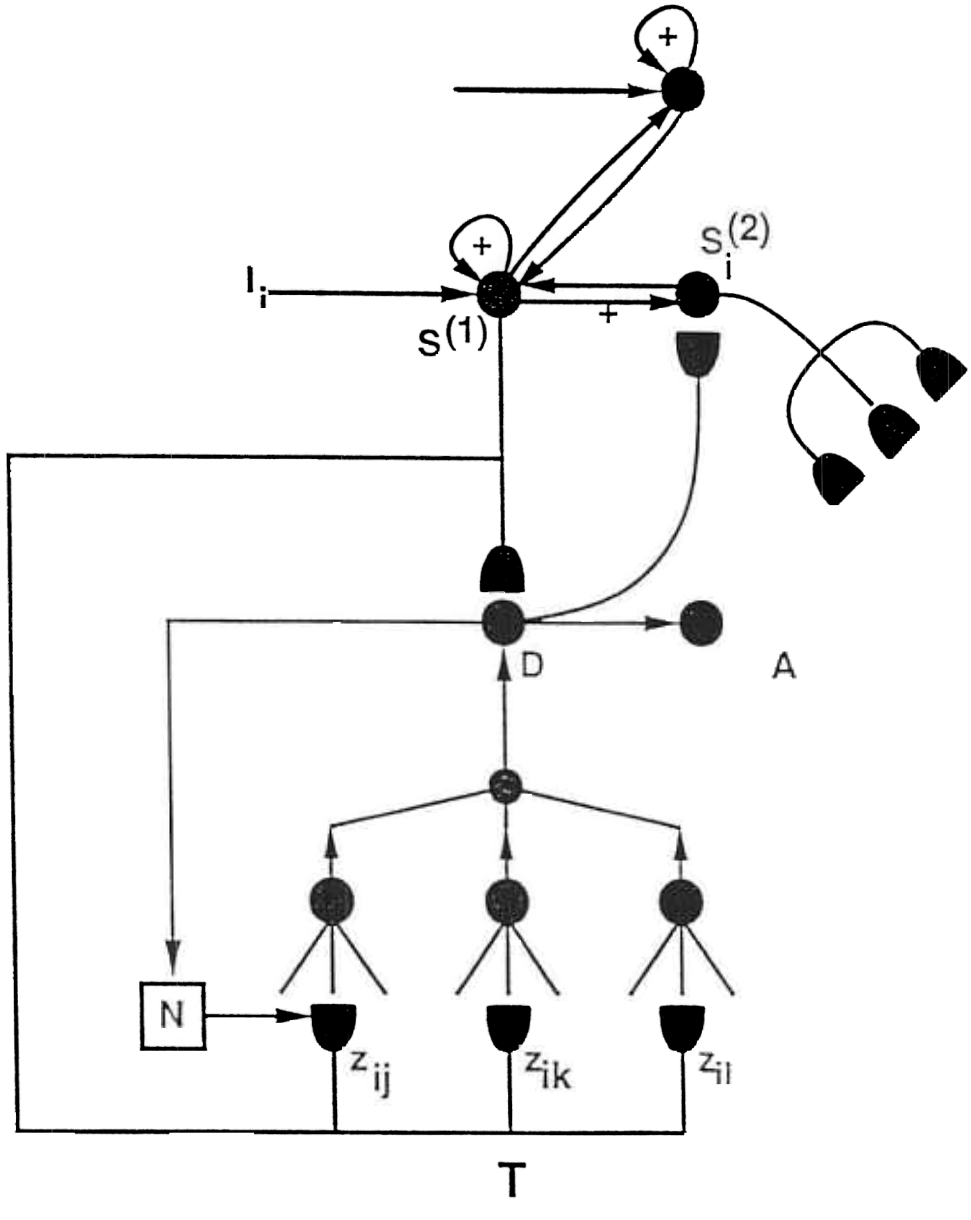

Figure

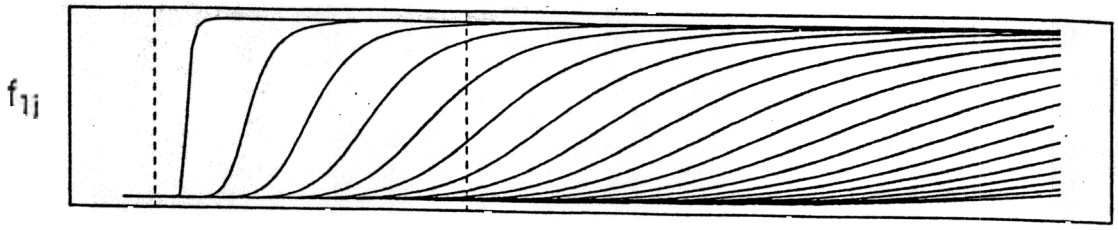

(a)

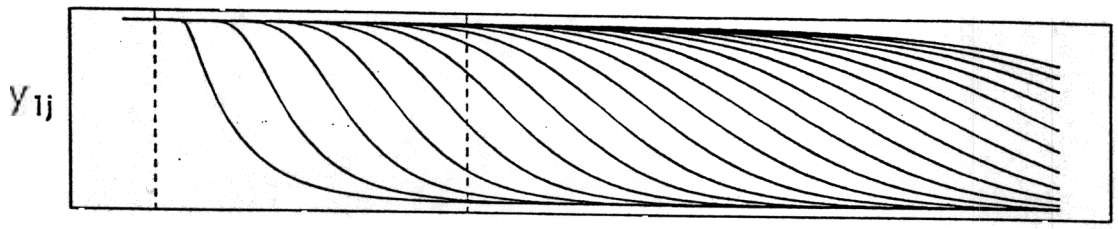

(b)

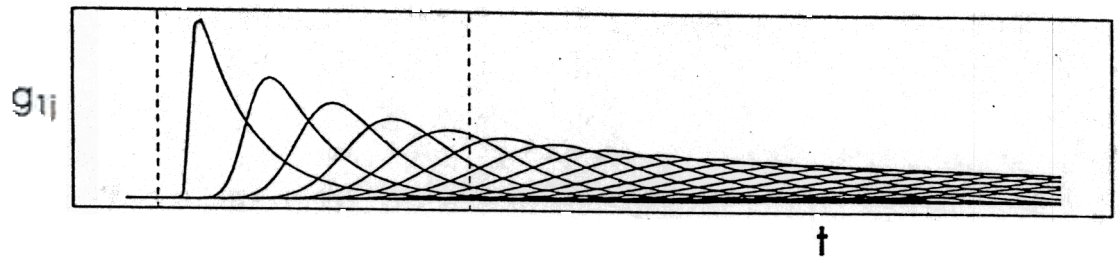

(c)

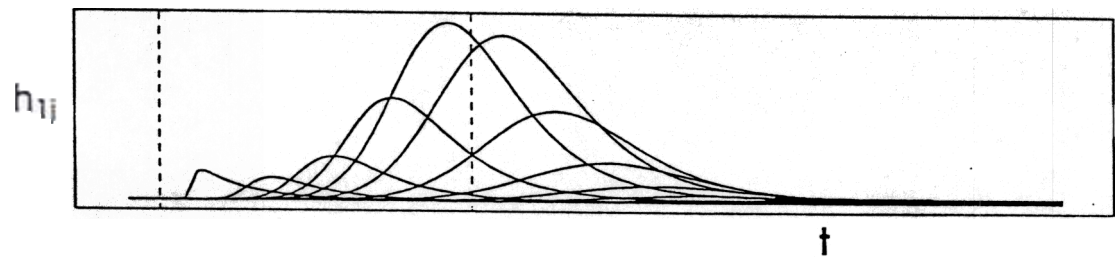

(d)

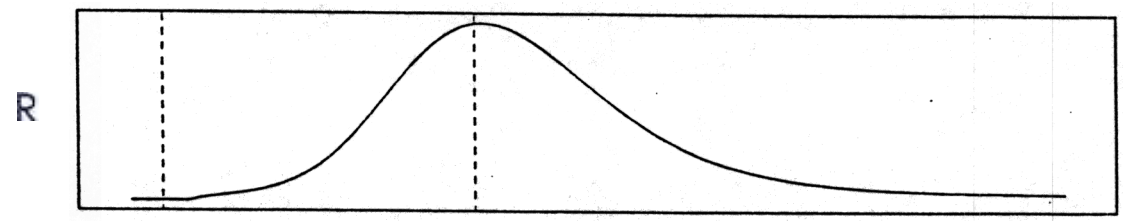

(e)

Figure

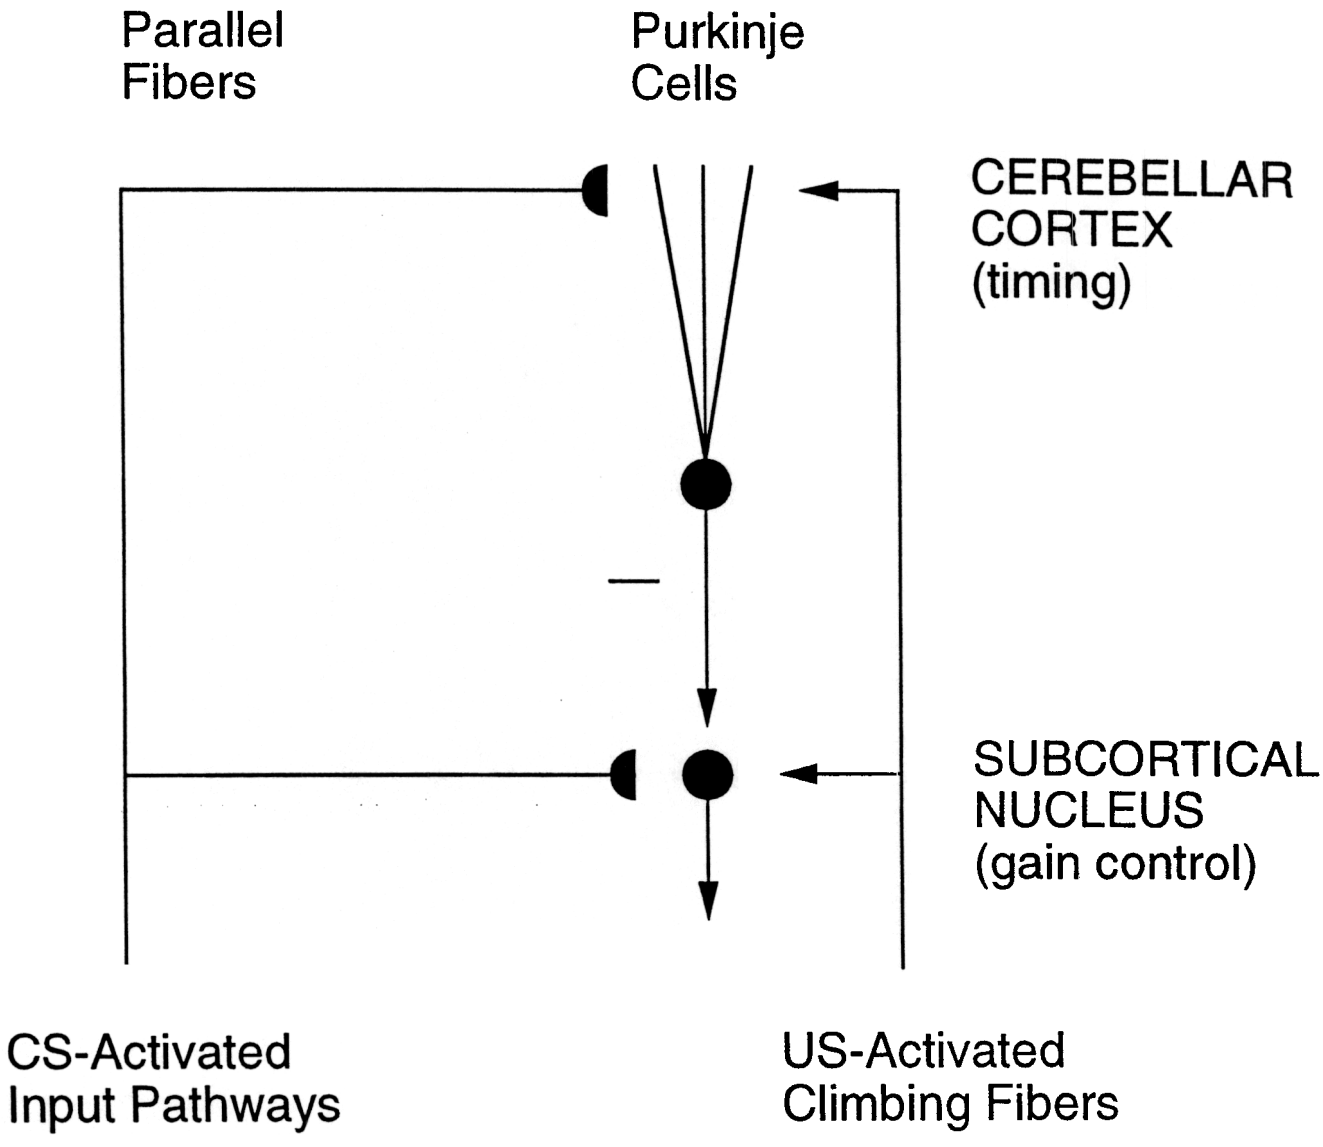
Parallel
Fibers
Purkinje
Cells
CEREBELLAR
CORTEX
(timing)
SUBCORTICAL
NUCLEUS
(gain control)
CS-Activated
Input Pathways
US-Activated
Climbing Fibers

Figure 6

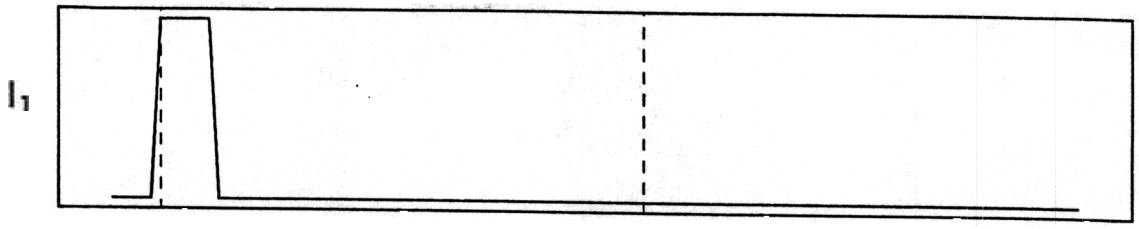

(a)

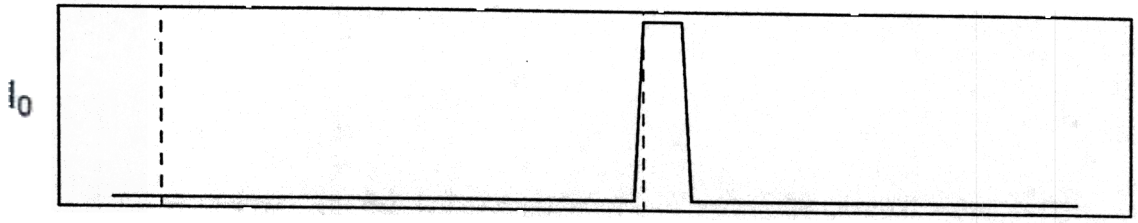

(b)

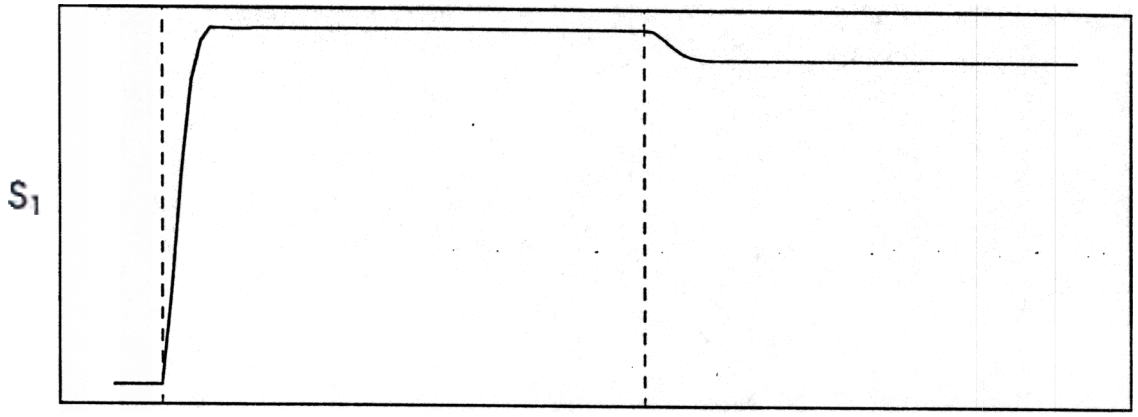

(c)

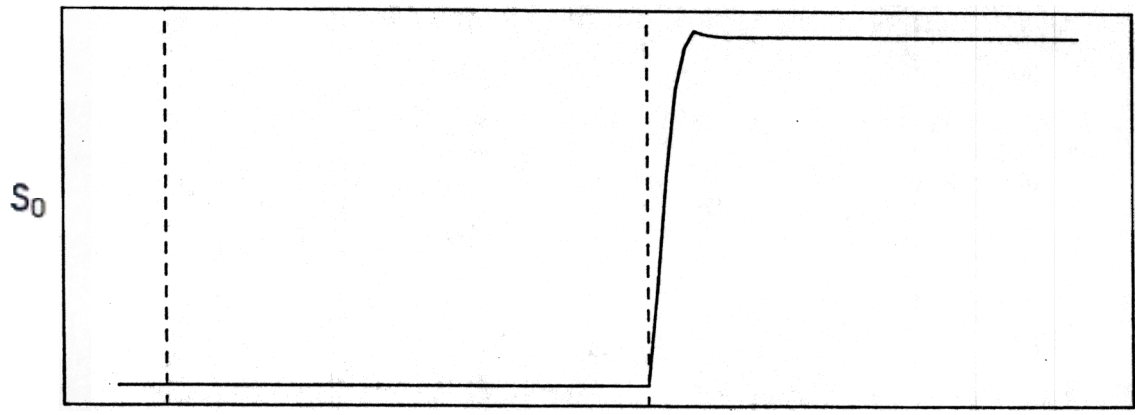

Figure

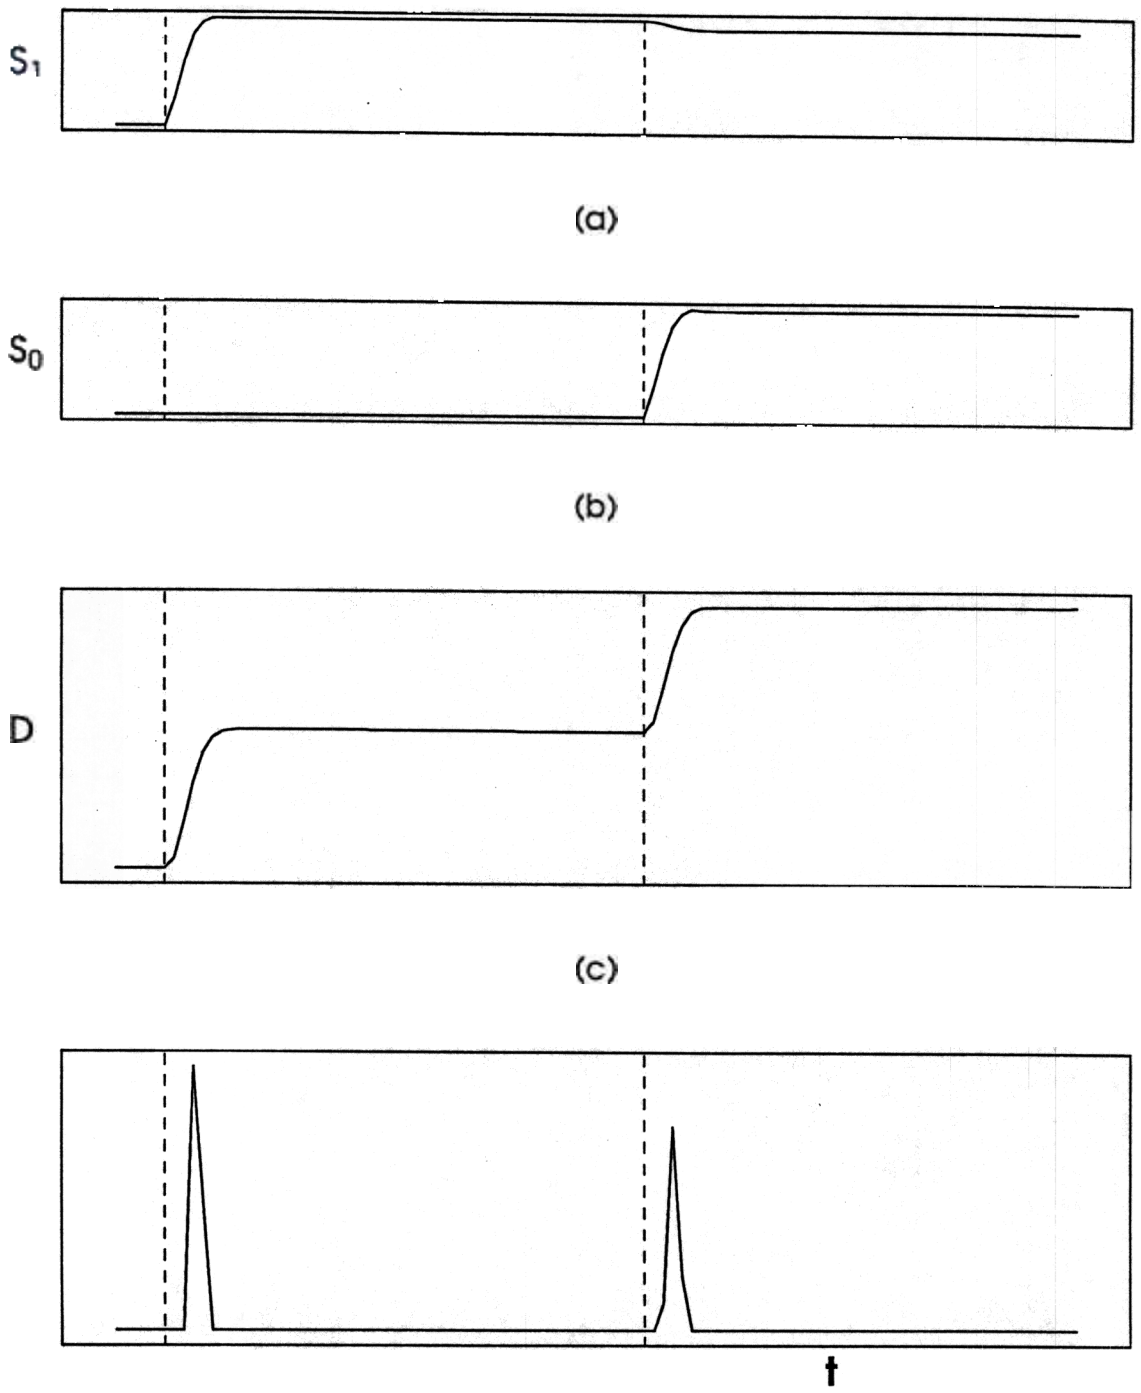

Figure

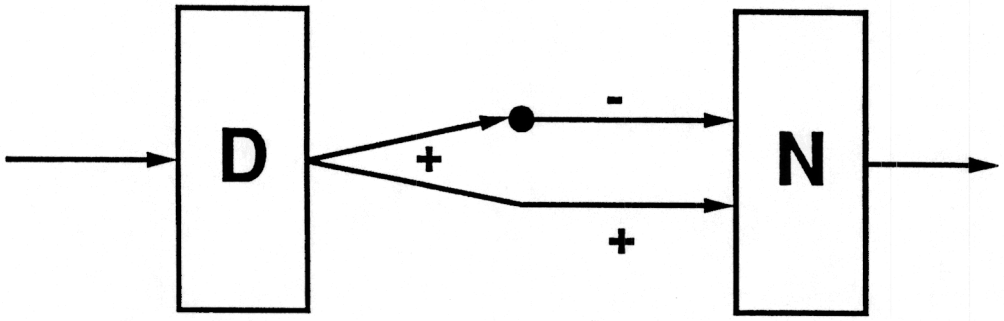

(a)

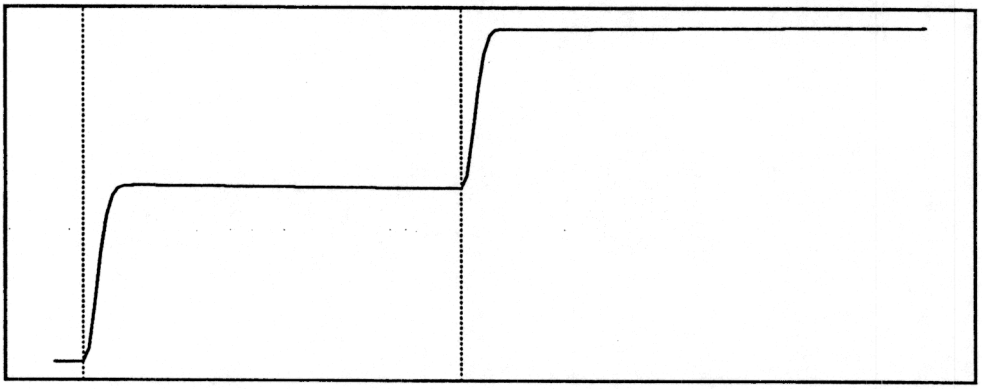

(b)

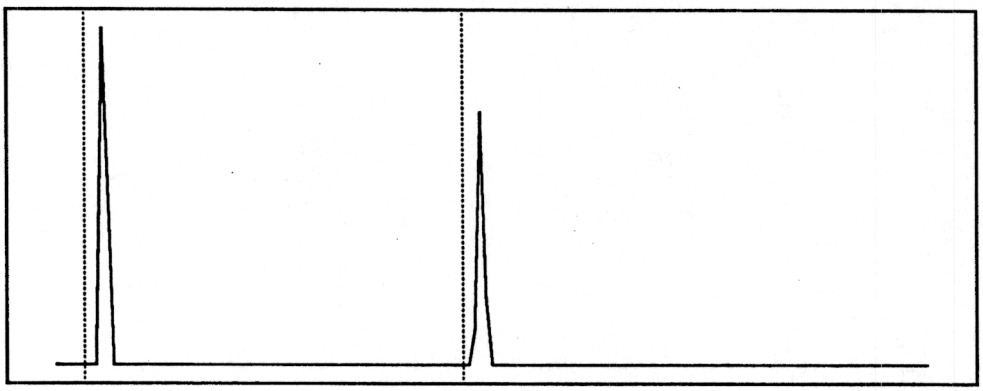

(c)

Figure

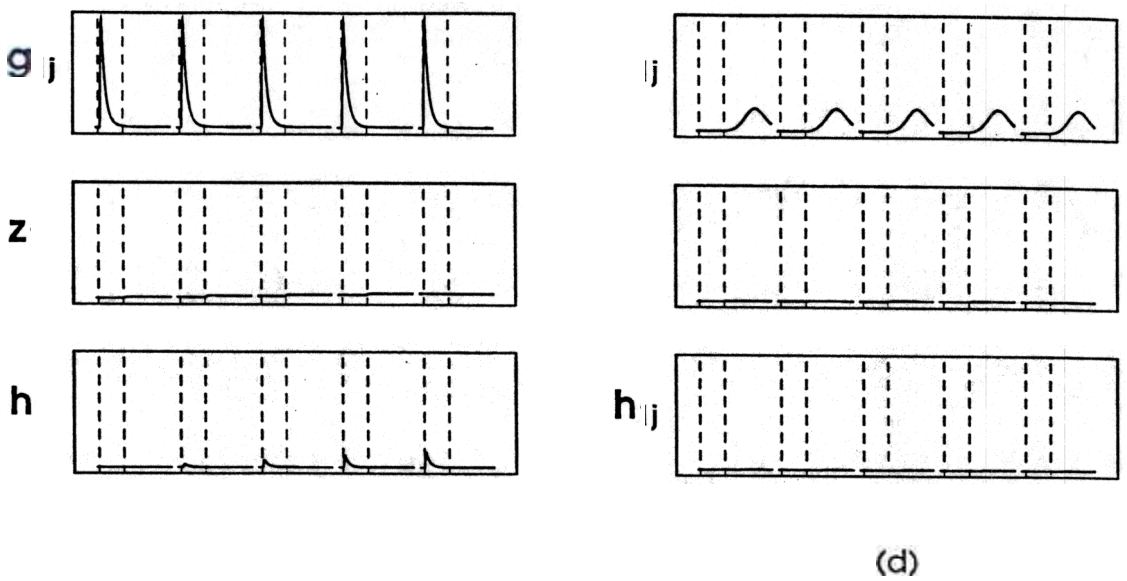

(d)

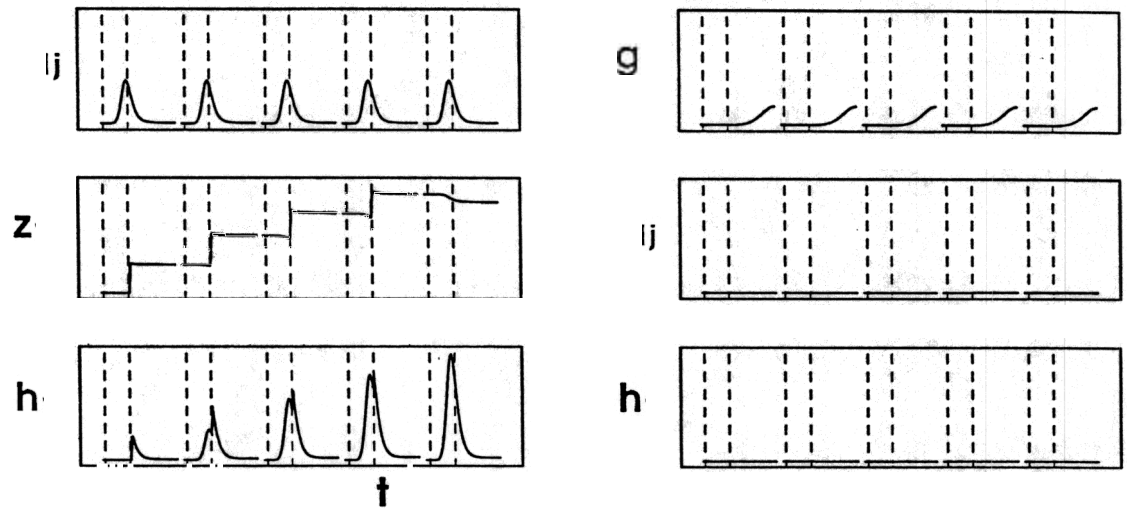

(e)

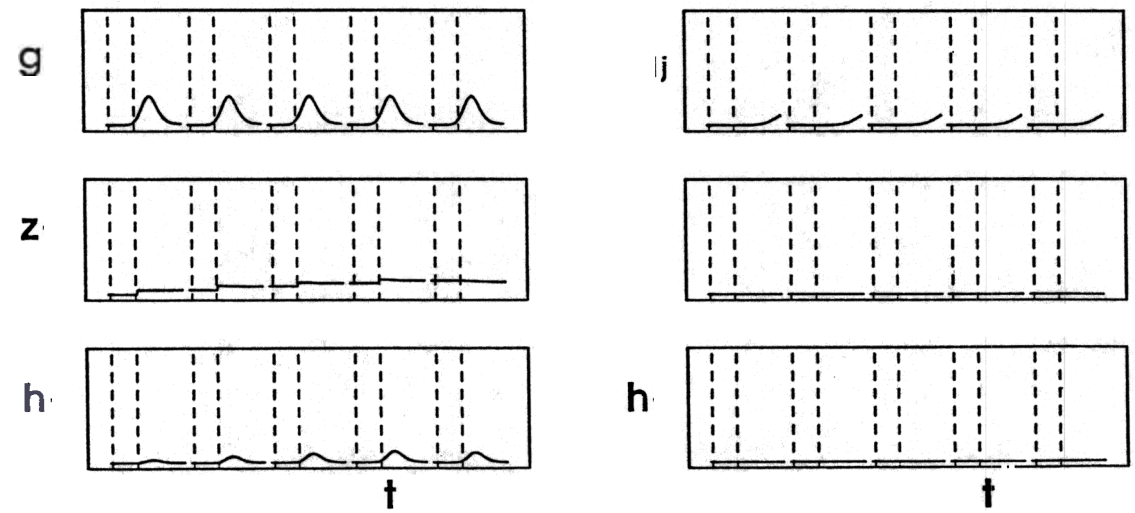

(f)

Fig

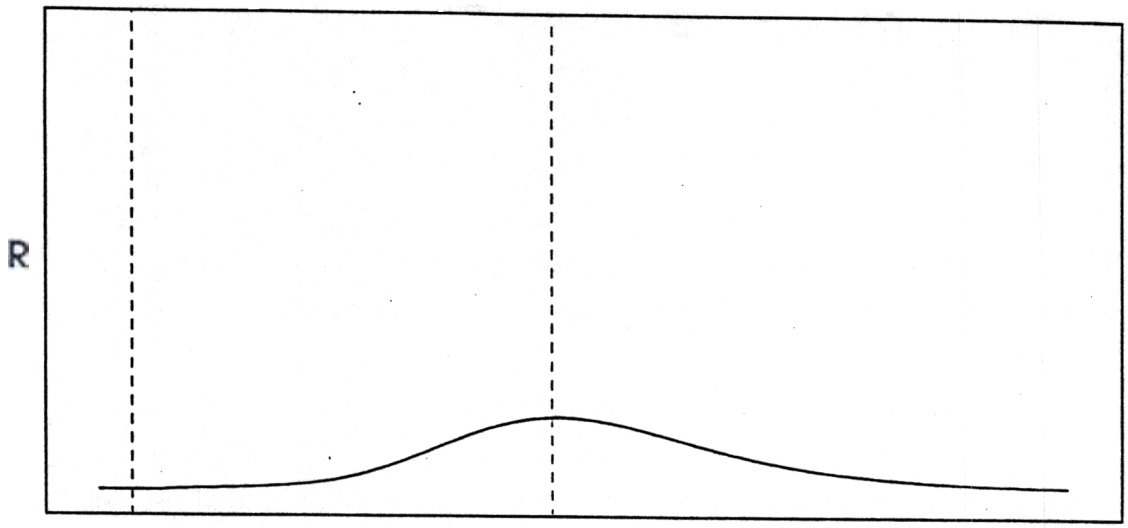

(a)

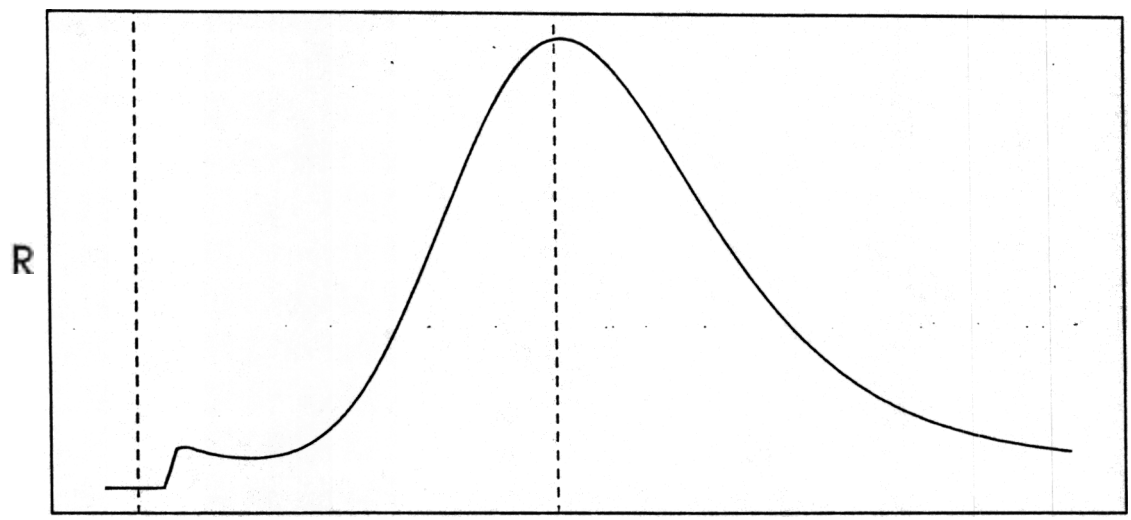

Figure

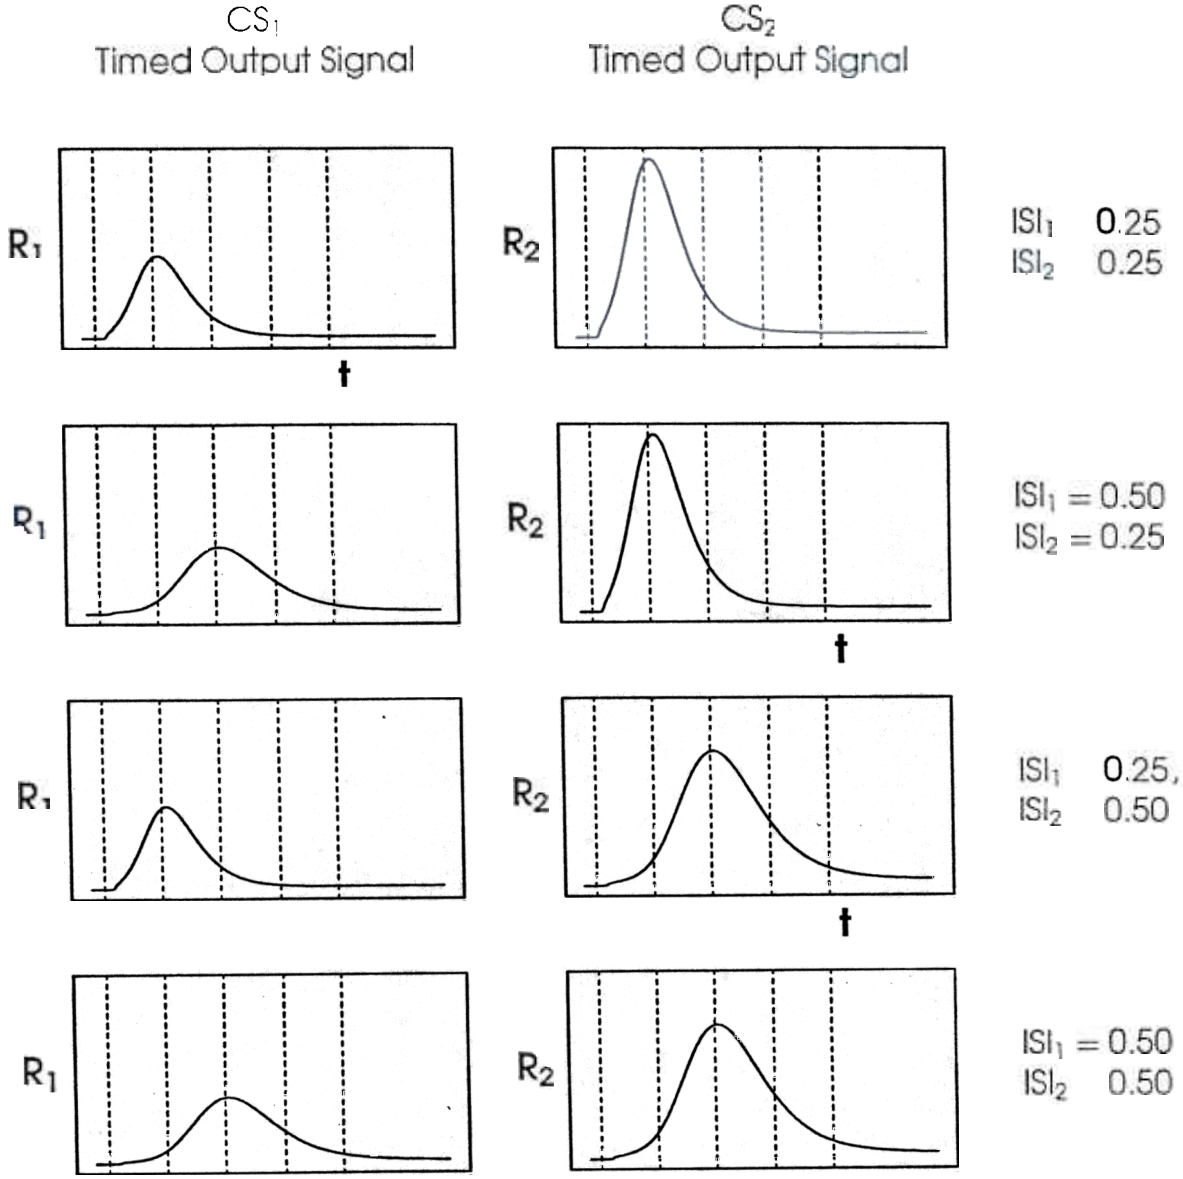
CS1
Timed Output Signal
CS2
Timed Output Signal
R1
R2
t
ISI1 0.25
ISI2 0.25
R1
R2
t
ISI1 = 0.50
ISI2 = 0.25
R1
R2
t
ISI1 0.25,
ISI2 0.50
R1
R2
ISI1 = 0.50
ISI2 0.50

Figure