

BRESC 30002

Review

A neural network model of adaptively timed reinforcement learning and hippocampal dynamics

Stephen Grossberg and John W.L. Merrill

Center for Adaptive Systems and Department of Cognitive and Neural Systems, Boston University, Boston, MA 02215 (USA)

(Accepted 17 March 1992)

Key words: Learning; Timing; Neural network; Reinforcement; Emotion; Recognition; Attention; Motor control; Hippocampus; Thalamus; Cortex; Cerebellum; *N*-Methyl-D-aspartate

CONTENTS

Abstract	4
1. Introduction and overview	5
1.1. Timing the balance between exploration for novel rewards and consummation of expected rewards	5
1.2. Distinguishing expected non-occurrences from unexpected non-occurrences: inhibiting the negative consequences of expected non-occurrences	6
 <i>Part I. Spectral timing</i>	
2. An example of spectral timing: conditioning the nictitating membrane response and the FI scallop	
3. The timing paradox and a solution	
4. START: a unified model of adaptive timing and conditioned reinforcer learning	
4.1. Limited capacity short term memory	
4.2. Drive representation	10
4.3. Conditioned reinforcement	11
4.4. Activation spectrum	11
4.5. Habituated transmitter spectrum	11
4.6. Gated signal spectrum	11
4.7. Spectral learning law	12
4.8. Now Print signal	12
4.9. Doubly gated signal spectrum	
4.10. Output signal	
 The problem of self-printing during adaptively timed secondary conditioning	
6. Simulations of secondary conditioning .	
7. The asymmetry between CS and US processing in timing control	
8. Different effects of US duration on emotional conditioning and adaptive timing: sustained and transient responses	16
9. The problem of combinatorial explosion: stimulus versus drive spectra	16
10. Stability of learning over many trials	17

11. Robustness of the model	17
12. Inverted U and Weber law	18
13. Comparison with data using multiple ISIs	18
14. ISI shift experiments	18
15. Partial reinforcement experiments	
16. Time averaging in response to multiple stimuli	19
17. Conditioning and the hippocampus	19
18. Adaptively timed conditioning of hippocampal pyramidal cells	20
19. Comparison of conditioned properties of hippocampal pyramids, NMDA receptors at dentate cells, and hippocampal afferents	20
20. Conditioning at dendritic spines of dentate granule cells	23
21. Convergence of dentate granule cells at CA3 pyramidal cells	23
22. NMDA receptors and adaptive timing	25
 <i>Part II. Reinforcement, recognition and motor learning</i>	
23. Reinforcement learning in vertebrates and invertebrates	25
24. Attentional focusing, blocking and the hippocampus	26
25. Subcortical fear conditioning and a cortical role in extinction	27
26. Conditioning in the cerebellum	28
27. Macrocircuit for sensory-cognitive processing: adaptive resonance theory	29
28. Hippocampal lesions and medial temporal amnesia	31
29. A synthesis of sensory-cognitive and cognitive-reinforcer circuits	32
30. Influences of hippocampectomy on conditioned timing	33
31. Concluding remarks: varieties of learning functions and networks	33
Acknowledgements	34
References	34
Appendix	38

A neural model is described of how adaptively timed reinforcement learning occurs. The adaptive timing circuit is suggested to exist in the hippocampus, and to involve convergence of dentate granule cells on CA3 pyramidal cells, and *N*-methyl-D-aspartate (NMDA) receptors. This circuit forms part of a model neural system for the coordinated control of recognition learning, reinforcement learning, and motor learning, whose properties clarify how an animal can learn to acquire a delayed reward. Behavioral and neural data are summarized in support of each processing stage of the system. The relevant anatomical sites are in thalamus, neocortex, hippocampus, hypothalamus, amygdala and cerebellum. Cerebellar influences on motor learning are distinguished from hippocampal influences on adaptive timing of reinforcement learning. The model simulates how damage to the hippocampal formation disrupts adaptive timing, eliminates attentional blocking and causes symptoms of medial temporal amnesia. Properties of learned expectations, attentional focussing, memory search and orienting reactions to novel events are used to analyze the blocking and amnesia data. The model also suggests how normal acquisition of subcortical emotional conditioning can occur after cortical ablation, even though extinction of emotional conditioning is retarded by cortical ablation. The model simulates how increasing the duration of an unconditioned stimulus increases the amplitude of emotional conditioning, but does not change adaptive timing; and how an increase in the intensity of a conditioned stimulus 'speeds up the clock', but an increase in the intensity of an unconditioned stimulus does not. Computer simulations of the model fit parametric conditioning data, including a Weber law property and an inverted U property. Both primary and secondary adaptively timed conditionings are simulated, as are data concerning conditioning using multiple interstimulus intervals (ISIs), gradually or abruptly changing ISIs, partial reinforcement and multiple stimuli that lead to time-averaging of responses. Neurobiologically testable predictions are made to facilitate further tests of the model.

1. INTRODUCTION AND OVERVIEW

This article contributes to the development of a behavioral and neurobiological theory of learning and memory. The theory describes how processes of learning, recognition, reinforcement and recall interact to focus attention upon motivationally desired goals, to generate appropriate actions towards these goals and to regulate selective forgetting of environmental contingencies that no longer predict behavioral success. Recent contributions to the theory are found in refs. 85, 91, 93, 94, 118, 119 and 149. Although derived from postulates aimed at explaining vertebrate behavior, the theory has also been applied to explain neurobiological data concerning classical conditioning of the invertebrate *Aplysia*²⁸. Other empirically supported predictions and as-yet-untested predictions of the theory are reviewed in refs. 36, 37 and 88. These models have also been incorporated into the control architecture of freely moving adaptive robots for use in technology^{5,6}.

The present article further develops a part of the theory, introduced in ref. 94, which analyzes how recognition events control motivated behaviors that are adaptively timed. Several different types of brain processes organize the temporal unfolding of serial order in behavior. The present model instantiates one type of timing control, called *spectral timing*, and shows how it can modulate the course of recognition learning, reinforcement learning, and the timed onset of a goal-oriented action. The model's formal processing stages are also compared with anatomical, neurophysiological and biochemical data about several brain regions, notably the hippocampal formation.

One of the main tasks of the present work is to model how processes such as adaptive timing, reinforcement learning, attention and motor learning differ, yet are linked in the control of behavior. Thus the exposition needs to describe several different types of circuits that form part of a larger neural system. These results were announced in ref. 92. Part I, Section 2, summarizes data concerning timed learning of the rabbit nictitating membrane response and the pigeon FI scallop. These data are used in Section 3 to suggest how the model solves a problem called the Timing Paradox. Section 4 describes the new spectral timing model, and illustrates its processes with computer simulations. Sections 5–16 describe how the model explains some difficult parametric conditioning data, notably data about secondary conditioning and the effects of changing stimulus intensity, frequency, duration and timing. These computer simulations show that the model can replicate quantitative properties of data from several types of conditioning experiments. All of

these computer simulations use a single set of parameters. Robustness of the model's properties is also demonstrated using different sets of parameters. Sections 17–22 interpret the adaptive timing mechanism in terms of interactions between dentate granule cells and CA3 pyramidal cells in the hippocampus, notably at NMDA receptors. Neurobiological data in support of this hypothesis are summarized and new predictions are made.

Part II, beginning with Section 23, shows how the spectral timing model may be embedded into a larger neural system for the control of recognition learning, reinforcement learning and motor learning. These sections also summarize behavioral and neural data in support of each processing stage of this model system. The relevant anatomical sites are in thalamus, neocortex, hippocampus, hypothalamus, amygdala and cerebellum. The behavioral data include an explanation of blocking in normal animals, elimination of blocking in hippocampectomized animals, impairing timing in hippocampectomized animals, medial temporal amnesia in hippocampectomized animals, subcortical fear conditioning and abnormal fear extinction in animals with cortical lesions, and disruption of motor learning by cerebellar lesions. Various of these data were reported after the corresponding model stages were published. Such data illustrate the predictive power of the theory. No claim is made that all neural processes are modelled in this system. Rather, the system is a lumped model that attempts to provide a minimal representation of the processes that are rate-limiting in explaining the targeted data bases. The present work may be viewed as a step in the progressive unclumping of the model to analyze ever finer neural processing stages.

1.1. Timing the balance between exploration for novel rewards and consummation of expected rewards

The spectral timing model clarifies the following type of behavioral competence, which is stated in an informal fashion to emphasize its familiarity in our daily lives. Data which support and refine this discussion are summarized throughout the article. Many goal objects may be delayed subsequent to the actions that elicit them, or the environmental events that signal their subsequent arrival. Humans and many animal species can learn to wait for the anticipated arrival of a delayed goal object, even though its time of occurrence can vary from situation to situation. Such behavioral timing is important in the lives of animals which can explore their environments for novel sources of gratification. On the one hand, if an animal could not inhibit its exploratory behavior, then it could starve to death by restlessly moving from place to place, unable to

remain in one place long enough to obtain food there. On the other hand, if an animal inhibited its exploratory behavior for too long while waiting for an expected source of food to materialize, then it could starve to death if food is not, after all, forthcoming.

Thus an animal's survival may depend on its ability to accurately time the delay of a goal object based upon its previous experiences in a given situation. Such an animal needs to balance between its exploratory behavior which may discover novel sources of reward, and its consummatory behavior which may acquire expected sources of reward. To effectively control this balance, the animal needs to be able to suppress its exploratory behavior and focus its attention upon an expected source of reward at around the time that the expected delay transpires for acquiring the reward.

Concepts of attention, expectation, exploration, timing, novelty and reward have a long history in psychology^{95,96,106,107,138,155,169}. The present article contributes to a mechanistic understanding of these concepts by showing how rigorously defined neural network models can be used to explain a range of behavioral and neural data that have not yet been explained either at all, or in a unified fashion, by alternative means. No less important than data analyses per se is the theoretical explication of new organizational principles which clarify the environmental and computational problems that the models are designed to solve.

1.2. Distinguishing expected non-occurrences from unexpected non-occurrences: inhibiting the negative consequences of expected non-occurrences

In this regard, an intuitive concept such as 'timing' is insufficient to characterize the several functionally distinct types of timing mechanisms that the brain uses to organize ongoing behaviors. Spectral timing, in particular, calibrates the delay of a single behavioral act, rather than the organization of a correctly timed and speed-controlled sequence of acts, or the types of timing whereby circadian, ultradian, or motor rhythms are organized. The types of task to which spectral timing contributes may be motivated by the following example, which we again describe in intuitive terms for vividness. Suppose that an animal typically receives food from a food magazine 2 s after pushing a lever, and that the animal orients to the food magazine right after pushing the lever. When the animal inspects the food magazine, it perceives the non-occurrence of food during the subsequent 2 s. These non-occurrences 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 2 s, it can compute

this sensory disconfirmation many times before the 2-s delay has elapsed.

The central issue is: what spares the animal from erroneously reacting to these *expected non-occurrences* of food during the first 2 s as predictive failures? Why does the animal not immediately become so frustrated by the non-occurrence 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 2 s have elapsed, why does the animal then react to the *unexpected non-occurrence* of food by becoming frustrated, shifting its attention, and releasing exploratory behavior?

Grossberg and Schmajuk⁹⁴ 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.

PART I. SPECTRAL TIMING

2. EXAMPLES OF SPECTRAL TIMING: CONDITIONING THE NICTITATING MEMBRANE RESPONSE AND THE FI SCALLOP

A well-studied example of spectral timing is the conditioning of the rabbit nictitating membrane response. Rabbits, like many mammals, have a translucent sheet of tissue called a 'nictitating membrane', that acts as a third eyelid. The nictitating membrane response, which extends this sheet across the eye, can be classically conditioned. For example, a conditioned stimulus, or CS, can be paired with noxious unconditioned stimulus, or US, such as a periorbital shock or airpuff, that elicits membrane extension. Smith¹⁵³ studied the effect of manipulating the time lag between CS onset and US onset. This lag is called the *interstimulus interval*, or ISI. The CS was a 50-ms tone and the US was a 50-ms electric shock. The ISI values were 125, 250, 500 and 1,000 ms. The fact that conditioning occurred at ISIs much larger than the CS duration implies that an internal trace of the CS is stored in short term memory (STM) subsequent to CS offset. Because an internal CS trace is needed to bridge the ISI gap between CS offset and US onset, such a

paradigm is called trace conditioning, to distinguish it from the delay conditioning paradigm wherein the CS and US overlap in time.

Smith¹⁵³ found that the conditioned response, measured as percentage of responses and response amplitude, was determined by both ISI and US intensity, whereas response onset rate and peak time were determined by the ISI essentially independently of US intensity. An increase in the mean of the peak response time correlated with an increase in the variance of the response curve, for each ISI (Fig. 1). Grossberg and Schmajuk⁹⁴ called this the *Weber law property* after its similarity to psychophysical ratio scales observed in vision and audition. This similarity had also been observed and commented upon by Gibbon⁶⁰ in a discussion of scalar timing theory.

Fig. 2 summarizes a computer simulation by Grossberg and Schmajuk⁹⁴ of the Weber law property. The data and computer simulation in Figs. 1 and 2 clarify why this type of timing is called *spectral* timing. These curves show how the peak behavioral and model responses scale with the ISI. Below it is shown how this property emerges as a population response from an ensemble of model cells that react to the CS and US

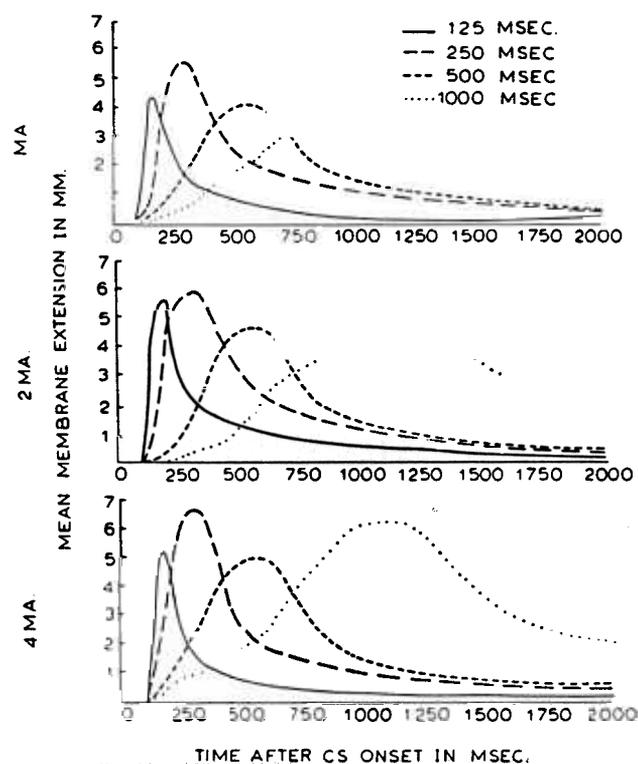


Fig. 1. Conditioning data from a nictitating membrane learning paradigm. Mean topography of nictitating membrane response after learning trial 10 with a 50-ms CS, ISIs of 125, 250, 500 and 1,000 ms, and different (1, 2, 4 MA) intensities of the shock US in each subsequent panel. (Reprinted from Smith¹⁵³ with permission.)

$$\sum_i f(x_i)y_i z_i$$

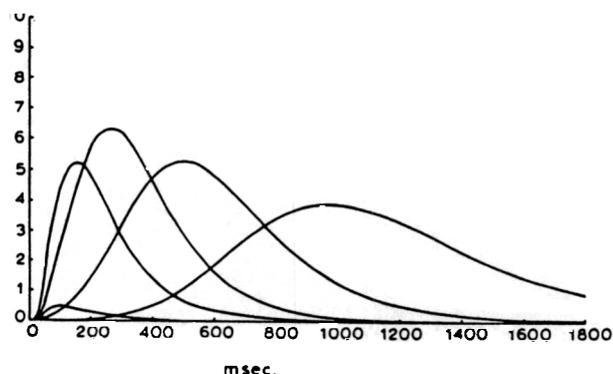


Fig. 2. Computer simulation of Weber law property and inverted U in learning as a function of ISI. The output signal functions $R(t) = \sum_i f(x_i)y_i z_i$ are plotted on a test trial, in response to the CS alone, subsequent to 10 prior learning trials with CS-US separated by different ISIs. Successive curves from left to right were generated by ISIs of 0 (the lowest amplitude curve), 125, 250, 500 and 1,000 ms using a US duration of 50 ms and an I_{US} intensity of 10 units. (Reprinted from Grossberg and Schmajuk⁹⁴ with permission.)

input at different rates. The set of rates defines a *spectrum* of activations that is densely distributed across all finite ISIs up to some maximum. The model shows how learning can enhance the activations of those cells whose response rates are best tuned to the experienced set of ISIs. Thus the ability of the model to scale so well with the behavioral response depends on its use of a cell population whose innate reaction rates may be sampled by learning to match the ISI.

Another example of spectral timing arises in appetitive instrumental conditioning with a fixed delay to reward in both rats^{127,144} and pigeons¹⁴⁵. In these experiments, animals were rewarded for the first lever press (in rats) or key peck (in pigeons) that occurred a fixed time after a signal was presented (Fig. 3). A characteristic pattern of response evolves: the animal withholds responding for a time, and then responds at an increasing rate. On test trials, the CS remains on for a long interval relative to the delay at which the US is normally presented, but no US occurs. Then the animal's response rate rises, before gradually falling. Examples of pigeon data are shown in Fig. 4.

These two different timed measures of response—nictitating membrane (NM) topography and fixed delay response rate—are obtained in different paradigms, yet they exhibit many common properties, including covariance of ISI with peak time and peak breadth. This kind of temporal covariation has been observed in many other paradigms, including signalled avoidance paradigms and differential reinforcement of long latency response (DRL) paradigms. These effects and others are reviewed in ref. 61.

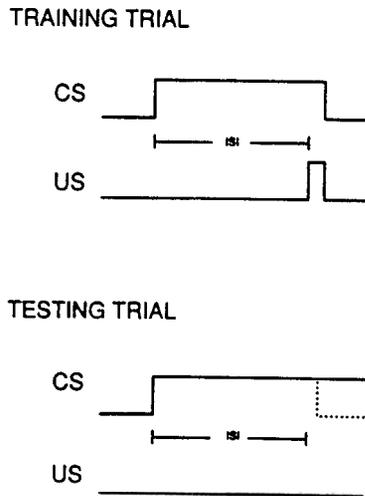


Fig. 3. A schematic of the training regimen in fixed-interval delay paradigm employed in the experiments summarized in Fig. 4. After the conditioned stimulus was turned on, the first key peck after a certain interval was rewarded. During testing trials, the conditioned stimulus was turned on and remained on for an interval far longer than the expected delay, so that the subjects' behavior at delays greater than the expected ISI could be quantified.

A variety of models may at the outset be imagined wherein the individual cells of a population differ by some parameter. The parameter could, for example, be the delay in a delay line through which each cell activates an output signal. Such a delay line model does not, however, easily generate the Weber law property or other data properties summarized and simulated in Sections 5–16.

3. THE TIMING PARADOX AND A SOLUTION

The Weber law property of a spectral timing model provides a way for an animal to distinguish between the expected and unexpected non-occurrences that were discussed in Section 1, without losing the capacity to time its conditioned responses. The Timing paradox described in this section clarifies why this is a non-trivial problem. The Timing paradox comprises the following, apparently contradictory, pair of constraints. On the one hand, in response to any fixed choice of conditionable ISI, the learned response delay approximates the ISI and thereby enables the animal to prepare appropriate responses for when they are most needed. Thus a model of adaptive timing needs to accurately discriminate between individual temporal delays. On the other hand, expected non-occurrences *throughout* the ISI should not be treated as predictive failures. Thus the inhibitory signal that prevents this from happening must be distributed throughout the ISI. How can a timing model both be sharply enough tuned to precisely learn the ISI, yet be broadly enough tuned to inhibit orienting responses throughout the entire ISI interval?

A spectral timing model reconciles the two requirements of accurate optimal temporal delay and temporally distributed activation via the Weber law property (Fig. 2). According to this property, the breadth of the model's temporal response scales with its peak time. Consequently the onset of the CS causes the immediate initiation of an output signal which is sustained

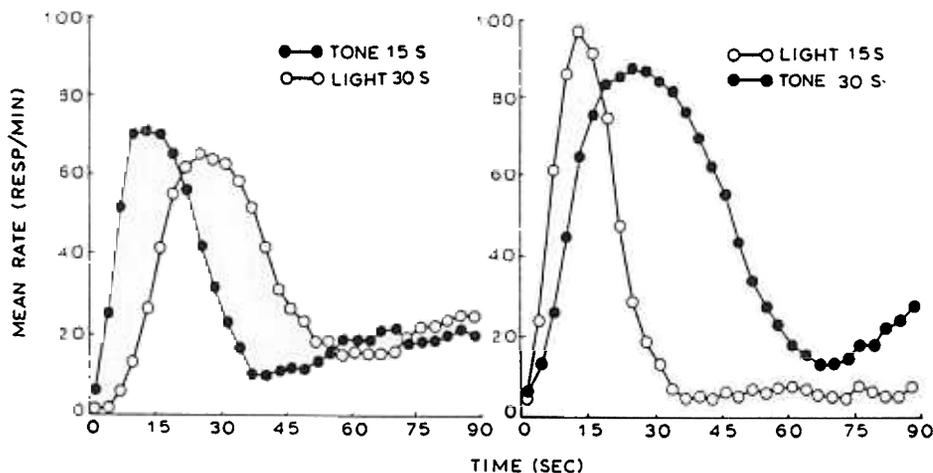


Fig. 4. Data on pigeon key pecking in a fixed-interval delay condition. All animals were trained to respond to two different conditioned stimuli, one visual and one auditory, each of which predicted reward for the first key peck after a fixed interval. For the animals in the group whose behavior is summarized in the graph on the left, the tone signalled availability of reward after a 15 second delay and the light signalled availability of reward after a 30-s delay. For animals in the group whose behavior is summarized in the graph on the right, the tone signalled availability of reward after a 30-s delay and the light signalled availability of reward after a 15-s delay. The times at which each response curve peaks correspond closely with the times at which each key peck is of maximal value. Also, within each stimulus modality, subjects' responses exhibit a covariation of peak time and peak breadth, as in the Weber law property shown in Figs. 1 and 2.

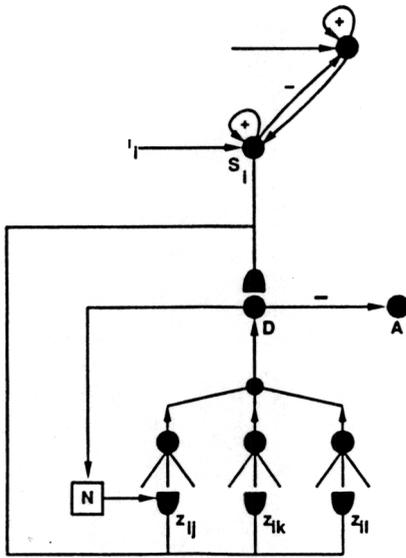


Fig. 5. A START model that combines a spectral timing module with a reinforcement learning network to achieve adaptively timed reinforcement learning and inhibition of the orienting subsystem.

throughout the entire ISI, but the *peak* output of the signal is accurately located at the expected arrival time of the US.

A solution of the Timing paradox is achieved by assuming that the output of the spectral timing model obeys the Weber law property, and that a neural circuit exists which processes this output to carry out two different functions along two different pathways. One pathway is used to inhibit the expression of a predictive failure, as shown in Fig. 5. By virtue of the Weber law property, this inhibition occurs throughout the ISI interval. The second pathway is used to excite, or energize, adaptively timed responding, as in Figs. 5 and 30 below. This property uses the fact that spectrally timed activations peak at the ISI. These pathways will be characterized with increasing precision in Section 4 and Part II.

4. START: A UNIFIED MODEL OF ADAPTIVE TIMING AND CONDITIONED REINFORCER LEARNING

The new adaptive timing model will now be defined. It combines Spectral Timing mechanisms with mechanisms from Adaptive Resonance Theory (see Part II). Hence it is called the START model. The START model builds upon a previous model of reinforcement learning whose processing stages are compared with behavioral and neural data below. Here we provide just enough exposition to define the model and to compare its emergent properties with these data.

As illustrated in Section 2, the model is tested by simulating data from reinforcement learning experiments, notably classical conditioning experiments. Each sensory event is therefore 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 (Fig. 5). 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 ⁷¹. 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 et al. have called them emotion nodes^{22,23} and Barto et al.⁸ 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 Fig. 5, both of these learning processes output to D , which in turn inhibits a population of cells called the orienting subsystem. The orienting subsystem is denoted by A because it is a source of non-specific arousal signals that are capable of initiating frustrative emotional reactions, attention shifts and orienting responses (see Part II). The inhibitory pathway from D to A is the gate that prevents these events from occurring in response to expected disconfirmations (Section 1).

4.1. 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^{75,81}. Expressed mathematically, each CS. acti-

vates an STM representation s_i whose activity S_i obeys the shunting on-center off-surround competitive feedback equation:

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

In (1), $I_i(t)$ 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 (1). Term $-\gamma_A S_i \sum_{k \neq i} f_S(S_k)$ describes inhibition of S_i by competitive signals $f_S(S_k)$ from the off-surround of populations $k \neq i$. Fig. 6 summarizes a computer simulation of how a brief CS_1 gives rise to a sustained

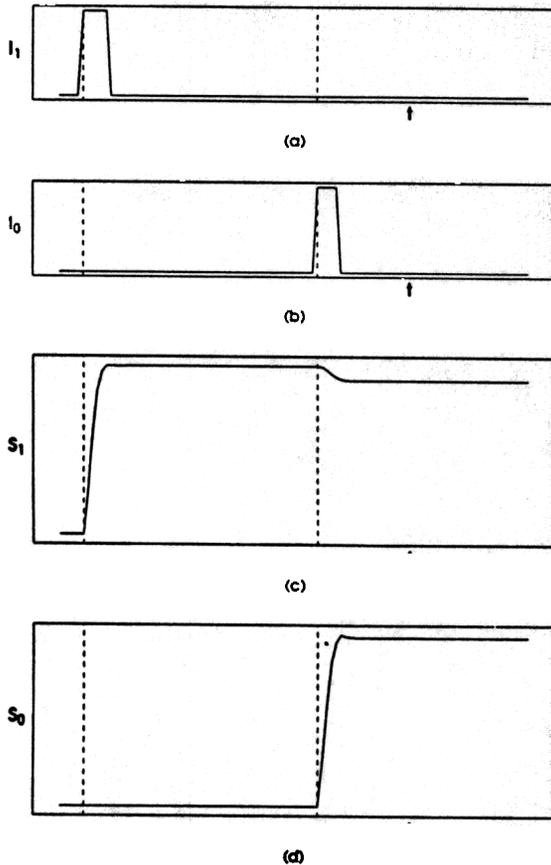


Fig. 6. In a START model, STM storage of a brief CS is achieved by positive feedback within the sensory representation s . CS attenuation by the US is dynamically controlled by the strength of recurrent inhibitory signals. a: input I_1 activated by CS_1 ; b: input I_0 activated by US; c: STM activation of CS_1 sensory representation; d: STM activation of US sensory representation.

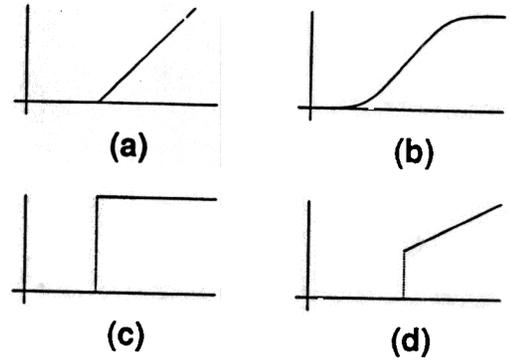


Fig. 7. Four possible feedback signal functions f for STM storage by equation (2): a: threshold-linear signal; b: sigmoid signal; c: binary signal; d: threshold-jump-linear signal.

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 forms depicted in Fig. 7 without qualitatively altering model properties. In this article, the simple rectification function

$$f(w) = [w - \mu]^+ \equiv \max(w - \mu, 0) \quad (2)$$

of Fig. 7a is used, except in equation (8) below, which uses a sigmoid signal function as in Fig. 7b.

4.2. 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^{83,85}. The activity D of the drive representation D obeys the equation

$$\frac{d}{dt} D = -\alpha_D D + \beta_D \sum_i f_D(S_i) C_i + \gamma_D R. \quad (3)$$

In (3), 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 (2), 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 (11).

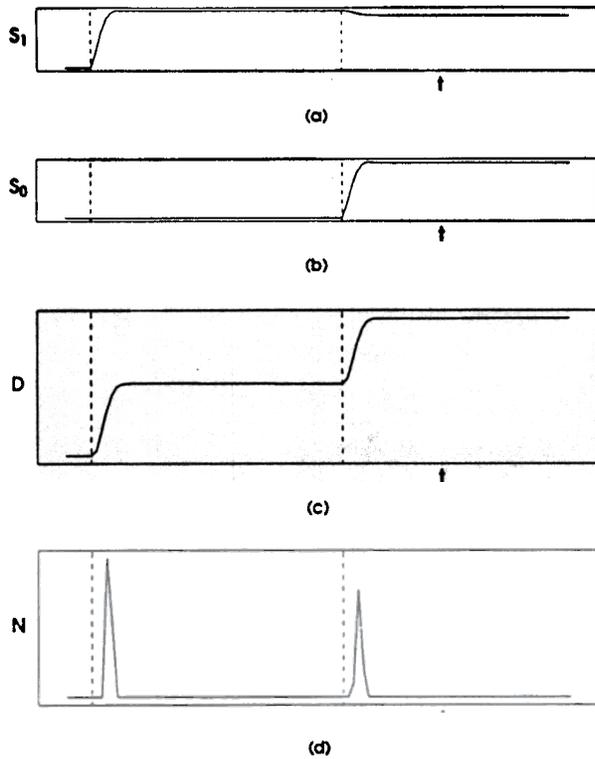


Fig. 8. Behavior of the Now Print module of the START model after many conditioning trials. a: activation of the sensory representation s_1 by the CS; b: activation of the sensory representation s_0 by the US; c: the resultant activation D of the drive representation D ; d: the resultant Now Print signal N .

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

4.3. Conditioned reinforcement

The adaptive weight C_i that calibrates conditioned reinforcement obeys a gated learning law⁶⁹:

$$\frac{d}{dt}C_i = \alpha_C S_i (-C_i + \beta_C (1 - C_i) f_C(D)). \quad (4)$$

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 (2). Activity C_i 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.

4.4. 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^{97,98}, and the spatial frequency-tuned cells of the visual cortex^{109,130,134-137,140,152,171,172,177}.

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

$$\frac{d}{dt}x_{ij} = r_j (-x_{ij} + (1 - x_{ij})f_x(S_i)), \quad (5)$$

where f_x satisfies (2). By (1) and (5), 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 , $j = 1, 2, \dots, n$. These potentials activate the next processing stage via signals

$$f(x_{ij}) = \frac{x_{ij}^8}{\delta_{ij}^8 + x_{ij}^8}. \quad (6)$$

Signal $f(x_{ij})$ is a sigmoid function of activity x_{ij} . Fig. 9a shows the activation spectrum $f(x_{ij}(t))$ that arises from presentation of CS_i to s_i via input I_i in (1), using a choice of rate parameters r_j in (5) which range from 10 (fast) to 0.0025 (slow). The method by which the simulations were performed is described in the Appendix.

4.5. Habituation transmitter spectrum

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

$$\frac{d}{dt}y_{ij} = \alpha_y (1 - y_{ij}) - \beta_y f(x_{ij}) y_{ij} \quad (7)$$

According to equation (7), 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 $y_{ij}(t)$, $j = 1, 2, \dots, n$, is called a habituation spectrum. The signal functions $f(x_{ij}(t))$ in Fig. 9a generate the habituation spectrum of $y_{ij}(t)$ curves in Fig. 9b.

4.6. 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

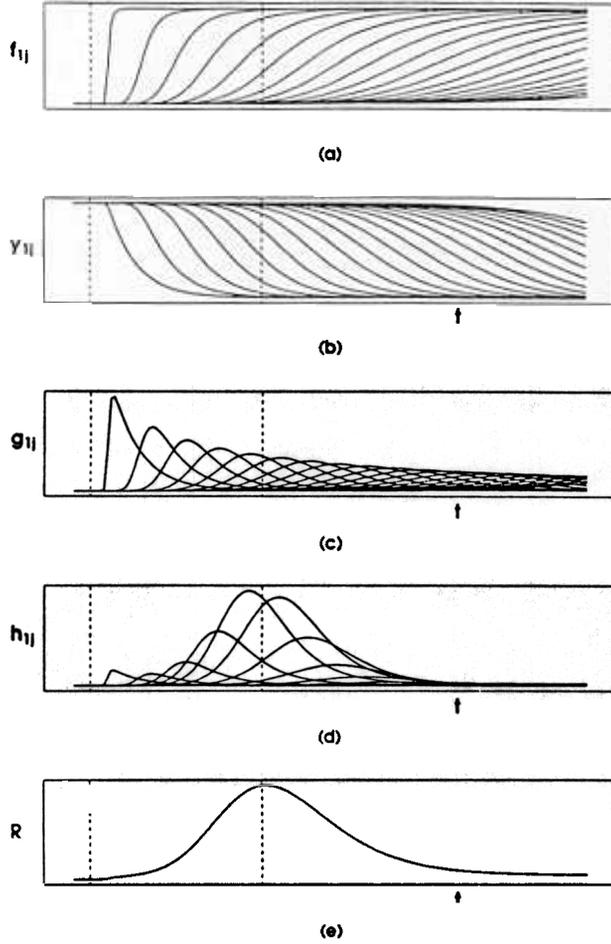


Fig. 9. Spectral timing properties of a START model. The CS₁ and US stimuli, were of intensity 10 for 0.05 time units and 0 otherwise. The onsets of CS₁ and US₁, respectively, are denoted in the figures by dotted vertical lines. The time difference between them of 0.5 time units is the ISI in the trial in question. a: the CS-activated spectrum $f_{ij}(t) = f(x_{ij}(t))$; b: the habituating transmitter gates $y_{ij}(t)$; c: the singly gated spectrum $g_{ij}(t) = f(x_{ij}(t))y_{ij}(t)$; d: the doubly gated spectrum $h_{ij}(t) = f(x_{ij}(t))y_{ij}(t)z_{ij}(t)$ arising after 5 trials; e: the corresponding output signal $R(t)$. Simulations were performed as described in the Appendix, with parameters and signal functions given by $\gamma = 0.2$, $\alpha_y = 1.0$, $\beta_y = 125.0$, $\alpha_z = 1.0$, $\delta = 0.0$, $\epsilon = 0.02$, $\alpha_E = 240.0$, $\alpha_A = 1.2$, $\beta_A = 120.0$, $\gamma_A = 12.0$, $\alpha_D = 120.0$, $\beta_D = 120.0$, $\gamma_D = 0.0$, $f_D(S) = [S - 0.05]^+$, $\alpha_C = 0.5$, $\beta_C = 25.0$, $f_C(D) = [D - 0.05]^+$, $f_A(A) = [A - 0.1]^+$, $F_X(A) = [A - 0.7]^+$, $r_j = 10.125 / (0.0125 + j)$; and the intensities of the CS and US inputs I_i in (1) equal 2.

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 Fig. 9c. 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 j , and M_{ij} is a decreasing function of j .

These laws for the dynamics of a chemical transmitter were described in Grossberg^{67,68}. They capture the simplest first-order properties of a number of known

transmitter regulating steps⁴⁶, such as transmitter production (term α_y), feedback inhibition by an intermediate or final stage of production on a former stage (term $-\alpha_y y_{ij}$), and mass action transmitter inactivation (term $-\beta_y f(x_j)y_{ij}$). 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.

4.7. Spectral learning law

Learning of spectral timing obeys a gated steepest descent equation

$$\frac{d}{dt} z_{ij} = \alpha_z f(x_{ij}) y_{ij} (-z_{ij} + N) \quad (8)$$

where N is the Now Print signal of (9). Each long term memory (LTM) trace z_{ij} in (8) 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 (8) is an example of learning by *gated steepest descent*. 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 (4), z_{ij} can either increase or decrease as a result of learning.

Associative learning by gated steepest descent was incorporated into neural network models in Grossberg⁶⁹ and is the learning law that was used to introduce adaptive resonance theory^{77,78}. An associative learning law of this form was subsequently used by Levy et al. to model their data on hippocampal LTP^{120,121}. Singer¹⁵¹ has also used such a law to model his experiments on adaptive tuning of visual cortical cells during the visual critical period. These experiments support the adaptive resonance theory predictions^{77,78} that both hippocampal LTP and feature detector tuning in visual cortex should obey such a learning law.

4.8. Now Print signal

A transiently active Now Print signal N modulates the learning process of (8). The signal N may be activated either by a US or by a CS that has already become a conditioned reinforcer. Both the US and a conditioned reinforcer CS can activate the drive repre-

sensation D , as shown in (3). We assume that the Now Print signal N is turned on by sufficiently large and rapid increments in the activity D of D . The transient signal N is derived from the sustained activity D by the action of a slow inhibitory interneuron (Fig. 10). The transformation from sustained activity D to transient activity N can be realized mathematically by the function

$$N = [f_C(D) - E - \epsilon]^+ \quad (9)$$

In (9), E is the activity of an inhibitory interneuron that time-averages $f_C(D)$, as in equation

$$\frac{d}{dt}E = \alpha_E(-E + f_C(D)), \quad (10)$$

before inhibiting the direct excitatory signal $f_C(D)$. Equation (9) means that $N = 0$ if $f_C(D) - E \leq \epsilon$, and $N = f_C(D) - E - \epsilon$ if $f_C(D) - E > \epsilon$. Figs. 8d and 10c illustrate how N responds to increments in D . An important property of N is that it increases in amplitude, but not significantly in duration, in response to larger inputs $f_C(D)$.

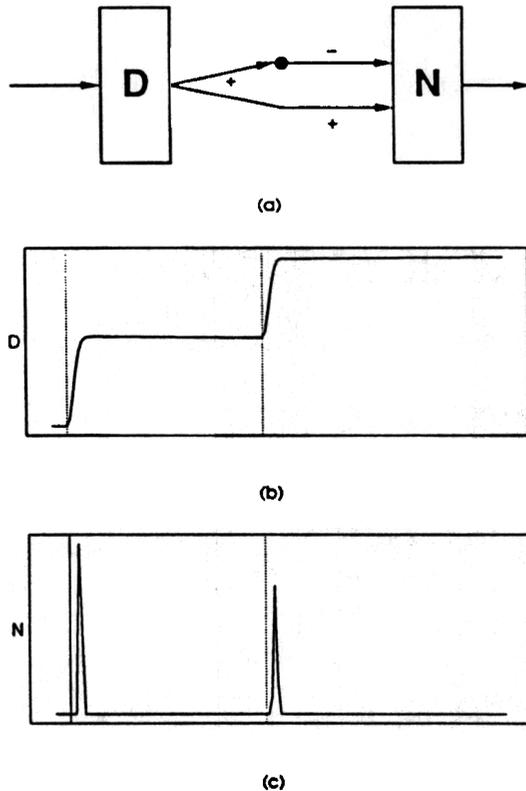


Fig. 10. Generation of a Now Print signal. a: the output of a drive representation D is converted into a Now Print signal N by passing this output through a fast excitatory pathway and a slower inhibitory pathway, whose signals converge at N . b: simulation of the activity D of D in response to two successive inputs, with the first response larger. c: activity N of N scales with the size of the increment in D . All parameters were as in Fig. 9.

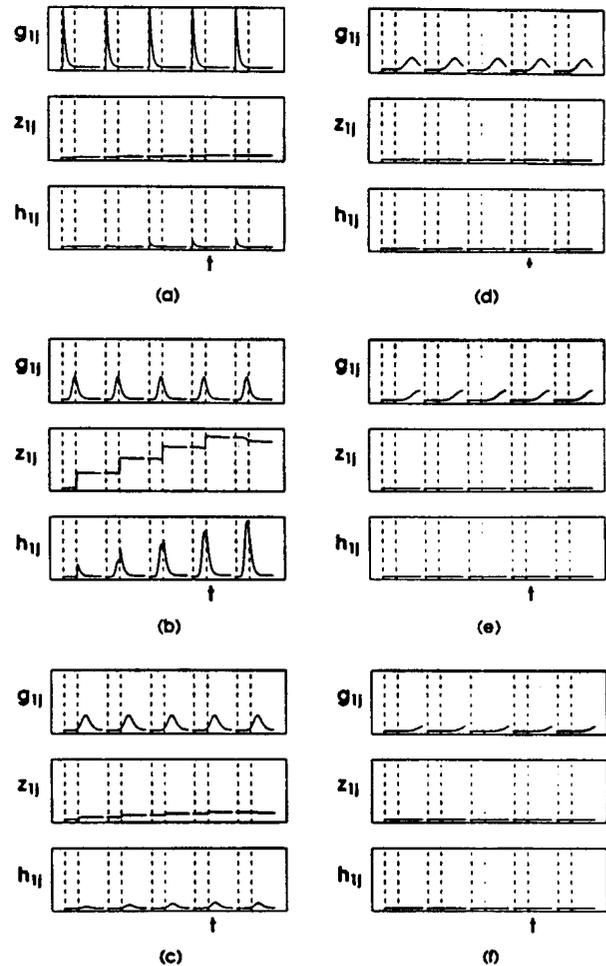


Fig. 11. Selective learning within different spectral populations at a fixed ISI = 0.5 time units. Each three-image panel from a to f represents the gated signal $g_{ij}(t)$ (top), long term memory trace $z_{ij}(t)$ (middle), and doubly gated signal $h_{ij}(t) = g_{ij}(t)z_{ij}(t)$ (bottom), at a different 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 Fig. 9 were used.

As noted above, the time interval between CS onset and US onset is called the interstimulus interval, or ISI. Using the spectral learning law (8)–(10), 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. Fig. 11 illustrates how 6 different LTM traces z_{ij} , $i = 1, \dots, 6$, 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 $z_{ij}(t)$.

4.9. 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 Fig. 11d, 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.

4.10. 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

$$R = \sum_{i,j} f(x_{ij})y_{ij}z_{ij}. \quad (11)$$

The output signal computes the cumulative learned reaction of the whole population to the input pattern. Fig. 9e shows the function R derived from the h_{ij} shown in Fig. 9d. A comparison of Fig. 9c–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 Fig. 9c are biased towards early times. The twice-gated signals $h_{ij}(t)$ in Fig. 9d are biased towards the ISI, but many signals peak at other times. The output $R(t)$ combines these partial views into a cumulative response that peaks at the ISI.

5. THE PROBLEM OF SELF-PRINTING DURING ADAPTIVELY TIMED SECONDARY CONDITIONING

The START model overcomes 4 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 described below, along with supportive data. Problems (1), (3) and (4) were not solved by the Grossberg and Schmajuk⁹⁴ 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^{63,112}. Indeed, Gormezano and Kehoe⁶³ 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 .

6. SIMULATIONS OF SECONDARY CONDITIONING

The START model overcomes the self-printing problem with its use of a transient Now Print signal N , as in (9). 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 (Fig. 10a,c). By the spectral learning law (8), negligible self-printing occurs. The main effect of the self-printing that does occur is to reduce every spectral LTM trace z_{ij} in (8) by a fixed proportion of its value, thus scaling down the size of $R(t)$ without changing the timing of its peak.

Fig. 12a depicts the model output $R(t)$ when the Now Print threshold ϵ in (9) 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. Fig. 12b shows the output when ϵ is set lower, thus allowing secondary conditioning and some self-printing to occur. Correct timing still obtains.

Fig. 13 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.

7. THE ASYMMETRY BETWEEN 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

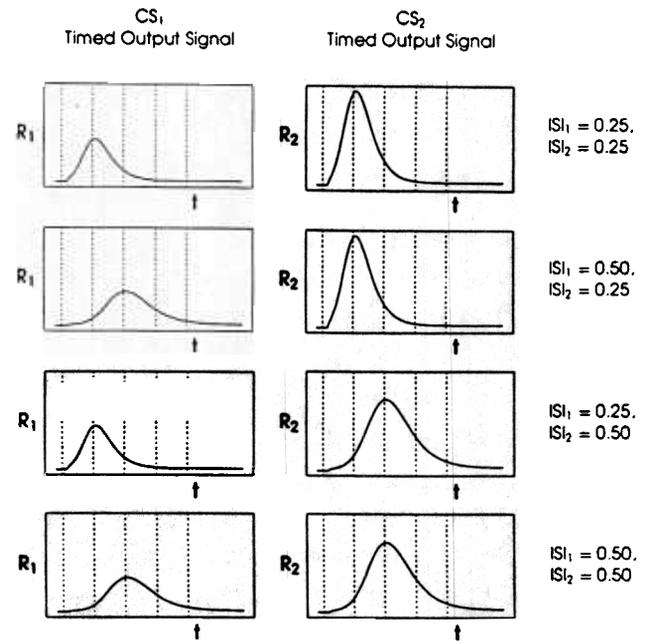


Fig. 13. START model output $R(t)$ during secondary conditioning with varying ISIs between the first and second CS_1 and between the second CS and the US using the parameters of Fig. 9. Notation ISI_1 below denotes the ISI between CS_1 and US, and ISI_2 denotes the ISI between CS_2 and CS_1 . On each 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, $ISI_1 = 0.25, ISI_2 = 0.25$; c, d, $ISI_1 = 0.5, ISI_2 = 0.25$; e, f, $ISI_1 = 0.25, ISI_2 = 0.5$; g, h, $ISI_1 = 0.5, ISI_2 = 0.5$.

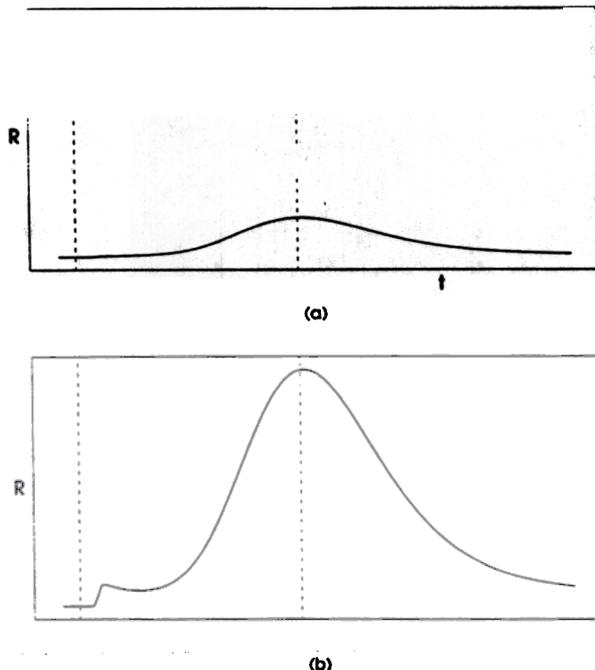


Fig. 12. The effect of self-printing upon the output of the model. a: a large threshold ϵ in the Now Print signal abolishes self-printing and secondary conditioning. It generates the lower output $R(t)$. b: a smaller threshold allows secondary conditioning and self-printing without a loss of timing. It generates the larger output $R(t)$.

reinforcer and an unconditioned stimulus are interchangeable. In fact, increasing the intensity of a conditioned reinforcer CS_1 can 'speed up the clock'^{123,127,176}, whereas increasing the intensity of a primary US can increase the amplitude of conditioned response, but does not change its timing¹⁵³.

The fact that parametric changes of CS and US may 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. Our proposed solution of the self-printing problem 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 (1). The larger input causes a larger peak amplitude of activity S_1 in (1), and a larger signal $f_x(S_1)$ in (5). By (5), 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 non-zero. Thus the total output R in (11) peaks at an earlier time, and causes the total output D from D in (3) 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 (5). This is true because a large US generates a signal $f_x(S_0)$ to the spectral activations in (5) at the same time that it generates a large signal $f_D(S_0)$ to D in (3) and a large Now Print signal N in (8). Thus a US creates the conditions of a 'zero ISI experiment' for purposes of spectral learning. All the LTM traces z_{0j} in (8) 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 (11), 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 (3) 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. See Grossberg and Schmajuk⁹⁴ for a further discussion of this issue.

8. 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^{4,16,19,40,111,164}. 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 (3) and (4); see ref. 73, Section 4 and ref. 81 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 Figs. 5 and 10. 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 in Sections 6–8 illustrate that the processes of emotional conditioning and adaptive timing, although linked, are not the same. They also support the START model's proposal of how these processes interact.

9. 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} , $j = 1, 2, \dots, n$. 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 z_{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 Fig. 5, each spectrum is

computed in parallel with its drive representation. Since the present simulations only consider 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. It is suggested in Section 20 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. Before then, computer simulations of the model which emulate data from various behavioral experiments will be summarized.

10. STABILITY OF LEARNING OVER MANY TRIALS

Some learning models become unstable when they experience a large number of learning trials. Fig. 14 shows the output of the model after 4, 50 and 100

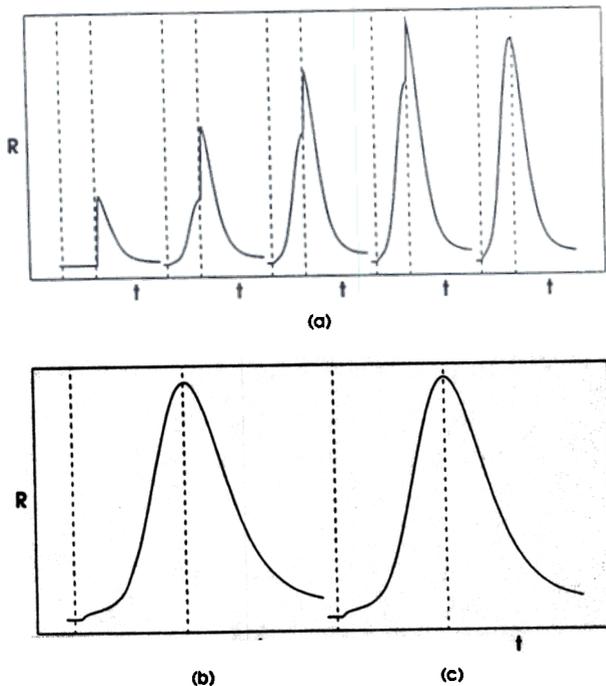


Fig. 14. Evolution of the model's output. Output $R(t)$ on each of the first 4 conditioning trials, followed by the CS alone on the fifth trial. Output after (b) 50 learning trials and (c) 100 learning trials.

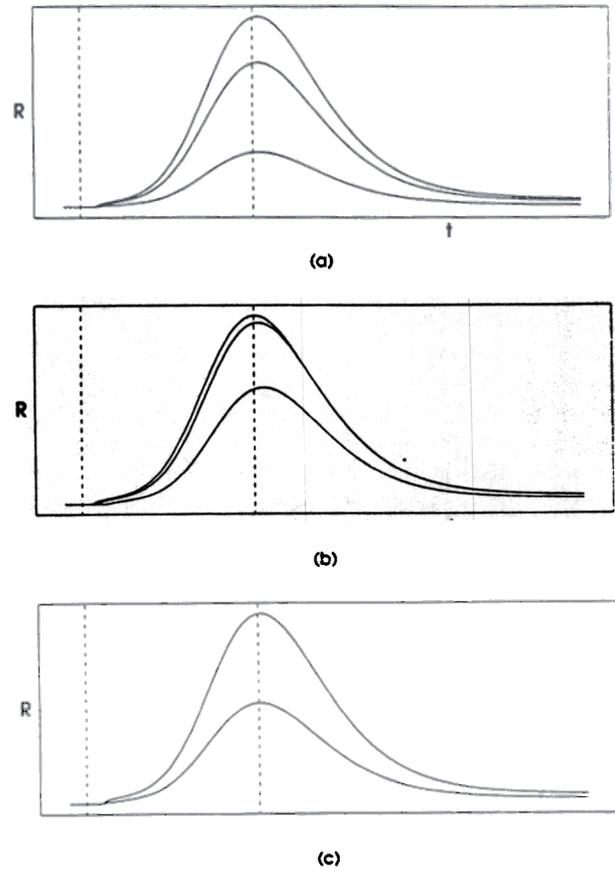


Fig. 15. Stability of learned timing under modifications of model parameters. a: effect of choosing the spectral learning rate α_z in (8) to $\frac{1}{2}$, 2 and 5 times that in Fig. 9. b: effect of setting the rates α_A , β_A and γ_A in (1) of the sensory representations to $\frac{1}{2}$, 1 and 2 times their values in Fig. 9. c: effect of proportionally changing the rates α_D , β_D , γ_D , and α_E in (3) and (10) to 2 and 4 times that in Fig. 16 in order to speed up the Now Print circuit.

learning trials, illustrating that this output pattern persists over many trials, even long after asymptote is reached, which occurs before trial 25 for the present choice of parameters.

11. ROBUSTNESS OF THE MODEL

Model properties are robust under physically plausible perturbations of its structure or parameters. For example, Fig. 15a shows that the asymptotic behavior of the model is qualitatively preserved under large changes in the learning rate α_z in (8). Fig. 15b shows that the model's behavior is unaffected by changes in the parameter which controls the speed at which the competition among sensory representations— s take place. Fig. 15c shows that the circuit's qualitative behavior is robust against large accelerations or decelerations of the rate at which D generates the Now Print signal.

The adaptive timing circuit learns accurately even when the behavior of some other part of the model is qualitatively altered. In Figs. 8 through 15, the parameters controlling the STM representations S were chosen so that STM can store more than one item. In Fig. 16, the parameters were chosen so that only one sensory representation can remain active through time. This has a dramatic effect upon the singly gated signals $f(x_{ij})y_{ij}$ within the model, since their support from S_1 does not persist when S_0 is large on training trials, but it has little effect upon the timing of the circuit, which again reaches maximal total activity R at around the time the US is expected.

12. INVERTED U AND WEBER LAW

Fig. 1 shows the average nictitating membrane (NM) topographies of animals trained with CS's of 50 ms duration and ISIs of 125, 250, 500, or 1,000 ms, depending upon the group to which the animal belonged.

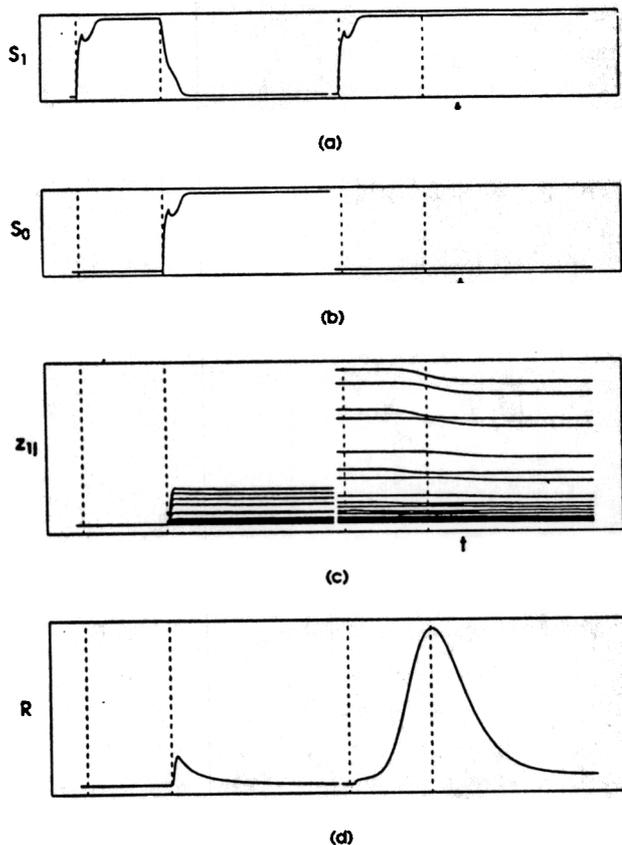


Fig. 16. Stability of learned timing when strengthened inputs to the sensory representation field no longer allows STM of the CS to remain active after the US is stored in STM: a: CS STM activity S_1 ; b: US STM activity S_0 ; c: spectral LTM traces after the first learning trial and in response to a CS alone on trial 25; d: output $R(t)$ under the conditions of c. STM parameters for this run: $\alpha_A = 0.6$, $\beta_A = 60.0$, $\gamma_A = 60.0$, US inputs I_1 and I_0 in (1) have intensities equal to 10.

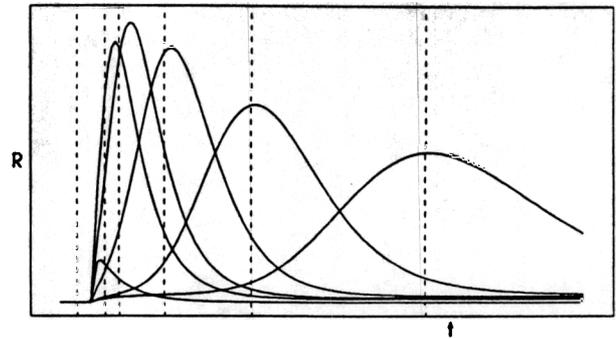


Fig. 17. An inverted U in output intensity as a function of CS-US ISI as produced by the START model with parameters as shown in Fig. 16. This figure was produced by plotting the outputs from the model to ISIs of 0.5, 0.125, 0.250, 0.500 and 1.000 units on a single set of axes.

Fig. 17 displays the outputs of the model at the corresponding ISIs. The model mimics the data pattern of rapid growth of the peak amplitude as the ISI increases through small values, followed by a gradual fall-off of peak amplitude for larger values of the ISI; and an increase in output width across time as the ISI increases.

13. COMPARISON WITH DATA USING MULTIPLE ISIS

Fig. 18 summarizes data from experiments, reported¹²⁸, in which rabbits were conditioned in a NM response paradigm. The ISI was one of two different values: 200 ms or 700 ms, with the different ISIs being presented at differing frequencies to different groups. In the group $P_{1/2}$ which received equal numbers of each ISI, the animals' average NM extension on test trials shows a double peak for the longer trials. The two peaks also exhibit the Weber law property. These double peak experiments provide strong evidence that a spectrum of possible times exists that is tuned by experience. Fig. 19 summarizes a computer simulation of that condition, which also exhibits two peaks that obey the Weber law at the two times at which the US would have been delivered.

14. ISI SHIFT EXPERIMENTS

In Coleman and Gormezano⁴³, animals were conditioned in a paradigm whose temporal characteristics were shifted either gradually or abruptly, from a 200-ms ISI to a 700-ms ISI, or conversely, during the course of the experiment. The animals' behaviors across learning trials are summarized in Fig. 20. Fig. 21 summarizes a set of computer simulations that qualitatively mimic the conditions of the original experiments.

15. PARTIAL REINFORCEMENT EXPERIMENTS

The classical conditioning circuit depicted in Fig. 5 forms part of a larger model neural system that is capable of explaining many data about operant conditioning (see refs. 81–83 and 85 and Part II for further discussion). Correspondingly, many operant conditioning data share similar properties with classical conditioning data. For example, the experiments of Roberts¹⁴⁴ used an operant rat lever-pressing task in which frequency of reinforcement was varied but the ISI was fixed. This manipulation altered the terminal level of responding to the stimulus, without changing the peak time of responding; that is, partial reinforcement affects the likelihood, but not the timing, of the response. The results are shown in Fig. 22.

A computer simulation of the same paradigm is shown in Fig. 23. As in the Roberts¹⁴⁴ data, only the level, not the time at which the output peaked, was affected by the probability of reinforcement.

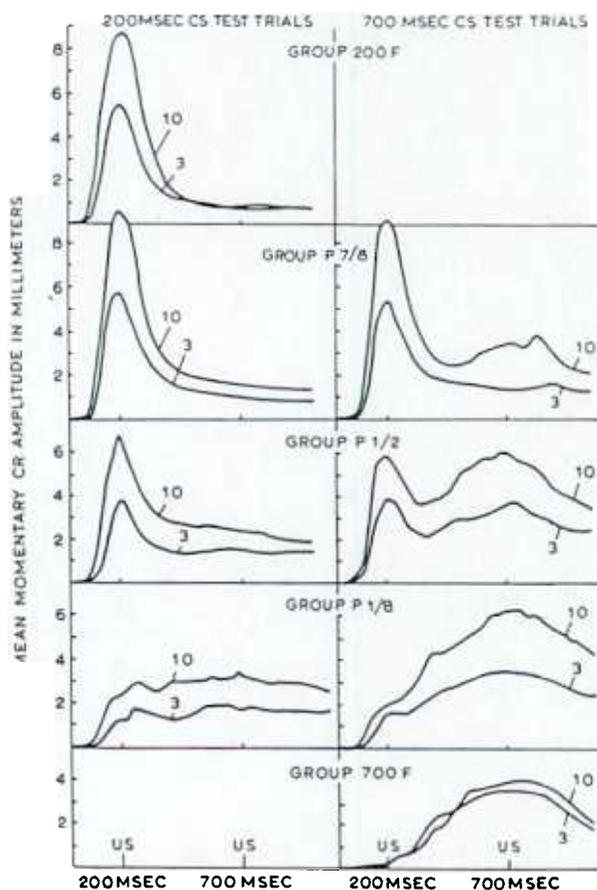


Fig. 18. Conditioning data from the nictitating membrane (NM) response paradigm as reported by Millenson et al.¹²⁸. Data shown are average NM extensions in CS-only trials with a tone CS of length 200 ms (left-hand panels) and 700 ms (right-hand panels) in animals presented with varying mixtures of training trials: 200-ms tone CSs immediately followed by a 50-ms shock US and 700-ms tone CSs immediately followed by a 50 ms shock US.

16. TIME AVERAGING IN RESPONSE TO MULTIPLE STIMULI

Holder and Roberts¹⁰¹ examined the effect of combining the timed responses to two different CS stimuli, a tone and a light, using rats in a lever-pressing task. If each of these stimuli has acquired the ability to elicit a conditioned response, and if they are presented sequentially, the resulting response is timed neither as the former nor as the latter stimulus would have required, but rather as an average. A simulation that qualitatively replicates this averaging property is shown in Fig. 24. This figure was generated with the model parameters set so that more than one sensory representation could be active in STM at one time. When the two stimuli were presented, the resulting output produced a peak that averages between the two expected times of arrival.

The comparisons between behavioral data and computer simulations in Figs. 8–24 illustrate how the spectral timing model emulates parametric behavioral properties from a number of conditioning paradigms. Sections 17–21 below point out that the formal model circuit also maps onto neural circuits in the hippocampus. This linkage provides a neural interpretation of anatomical and neurophysiological data concerning the role of the hippocampal formation in the control of adaptively timed conditioning. Testable predictions are made to further challenge this proposal. A brief historical discussion will first be given to clarify the larger neural modelling context in which the proposal needs to be evaluated.

17. CONDITIONING AND THE HIPPOCAMPUS

Learning within the $s \rightarrow D$ pathways of Fig. 5 was predicted in Grossberg^{71,76} to have the hippocampal formation as a final common path. It was also predicted that this type of learning is a variant of conditioned reinforcer learning. The distinction between the different learning processes that govern emotional conditioning and adaptive timing was not, however, made in these early articles. In experiments on conditioning the rabbit NM response, Berger and Thompson¹² reported that hippocampal learning does occur, thereby providing partial support for the prediction. At first, these investigators interpreted their results as the discovery of a general neural 'engram'. Subsequent experiments studied the effects of selective ablations on learning in both hippocampus and cerebellum¹²⁶. These experiments led to the conclusion that hippocampal learning appears to be a variant of the predicted condi-

tioned reinforcer learning, whereas the cerebellum carries out a type of motor learning. Thompson et al. (ref. 168, p. 82) distinguished these two types of learning as 'conditioned fear' and 'learning of the discrete adaptive response', respectively, a distinction that had also been predicted, and that is elaborated in Part II.

It should also be emphasized that this interpretation of hippocampal function does not contradict other data which implicate the hippocampal formation in the learning of spatial and attentional tasks^{103,131}. Such a hybrid functional role for hippocampus in conditioned reinforcement, spatial approach and avoidance, and attentional blocking was mechanistically outlined and predicted by the theory's earliest model circuits^{71,76} (reviewed in ref. 85). In support of such a hybrid function, Eichenbaum and Cohen⁵⁴ have summarized recent data showing that the same hippocampal cells which have place fields in a radial-arm maze can also show conditioned responses in classical conditioning tasks. The present article focuses on clarifying how the emotional conditioning and adaptive timing processes are designed and related to each other. In particular, as indicated in Fig. 5, reinforcement learning and adaptive timing are suggested to take place in different neural circuits, but circuits that interact with and modulate each other during normal behaviors. As reviewed below, aspects of emotional conditioning may be spared even if adaptive timing is deranged, just as aspects of motor performance may be spared even if adaptive timing is deranged.

18. ADAPTIVELY TIMED CONDITIONING OF HIPPOCAMPAL PYRAMIDAL CELLS

A large number of experiments have by now documented a role for hippocampal cells in adaptively timed conditioning. As Berger et al. (ref. 9, p. 204) have noted, "One of the striking features of these conditioning-induced changes in hippocampal activity is that a close parallel develops between the pattern of CS-evoked hippocampal pyramidal cell activity and the shape of the conditioned response—both during NM (nictitating membrane) aversive [10] and CJM (jaw movement) appetitive [14] paradigms." In addition, when animals were conditioned using different CS-US ISI intervals, the poststimulus histograms of pyramidal cell firing paralleled the shape of the NM response at ISIs of 150 ms and 250 ms. A 50-ms ISI did not lead to a conditioned NM response, and no enhancement of hippocampal unit activity occurred in either the CS-US interval or the US interval¹⁰⁰. In a signal detection task in which a white noise CS was varied from suprathresh-

old to threshold intensity, hippocampal firing to the CS completely predicted the occurrence of a behavioral response¹¹³. Finally, during an ISI shift experiment (see Section 14), Hoehler and Thompson¹⁰⁰ found that the peak time of the hippocampal trace changed before the peak time of the NM response topography. This difference may be analyzed in terms of hippocampal and cerebellar contributions to adaptive timing (see Section 26).

Such data led Berger, Thompson, and their colleagues to characterize the response pattern of hippocampal pyramidal cells as a 'temporal model' of the conditioned response, a proposal that was also espoused by Solomon^{157,158}. The START model suggests how this 'temporal model' develops and how it is integrated into a larger neural system for reinforcement learning, recognition learning and motor learning.

19. COMPARISON OF CONDITIONED PROPERTIES OF HIPPOCAMPAL PYRAMIDS, NMDA RECEPTORS AT DENTATE CELLS AND HIPPOCAMPAL AFFERENTS

Berger et al.⁹ reported data from dentate granule cells showing "increased firing rate beginning in the CS period and continuing through the US period... For any given cell, the latency of increased firing was constant and was time-locked to the CS" (p. 213). This difference between the 'time-locked' responses of dentate granule cells and the adaptively timed responses of hippocampal pyramidal cells suggests that pyramidal cells and dentate cells process hippocampal afferents in different ways. Berger et al.⁹ also reviewed data indicating that the high correlation between firing of hippocampal pyramidal cells and conditioned responses cannot be explained solely by conditioned changes in afferents to the hippocampus. In particular,

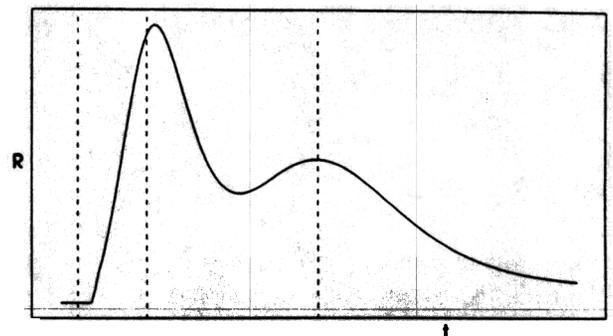


Fig. 19. Output $R(t)$ of the START model in a mixed ISI condition, such as in Fig. 18 (Panel $P_{1/2}$) in which 50% of all ISIs were 0.2 units and 50% of all ISIs were 0.7 units.

Berger et al.^{9,13} reviewed data about firing patterns in two major hippocampal afferents, the medial septum and the entorhinal cortex, during conditioning of the NM response. In the medial septum, each new stimulus generates a short, transient burst of activity followed by rapid habituation to the baseline response

pattern. The firing rate pattern in the entorhinal cortex is more like that seen in areas CA3 and CA1 of the hippocampus. However, the hippocampal behavioral trace is much stronger than the entorhinal trace and evolves more slowly. Whereas the entorhinal trace takes only 10–20 trials to form and to reach its asymptotic

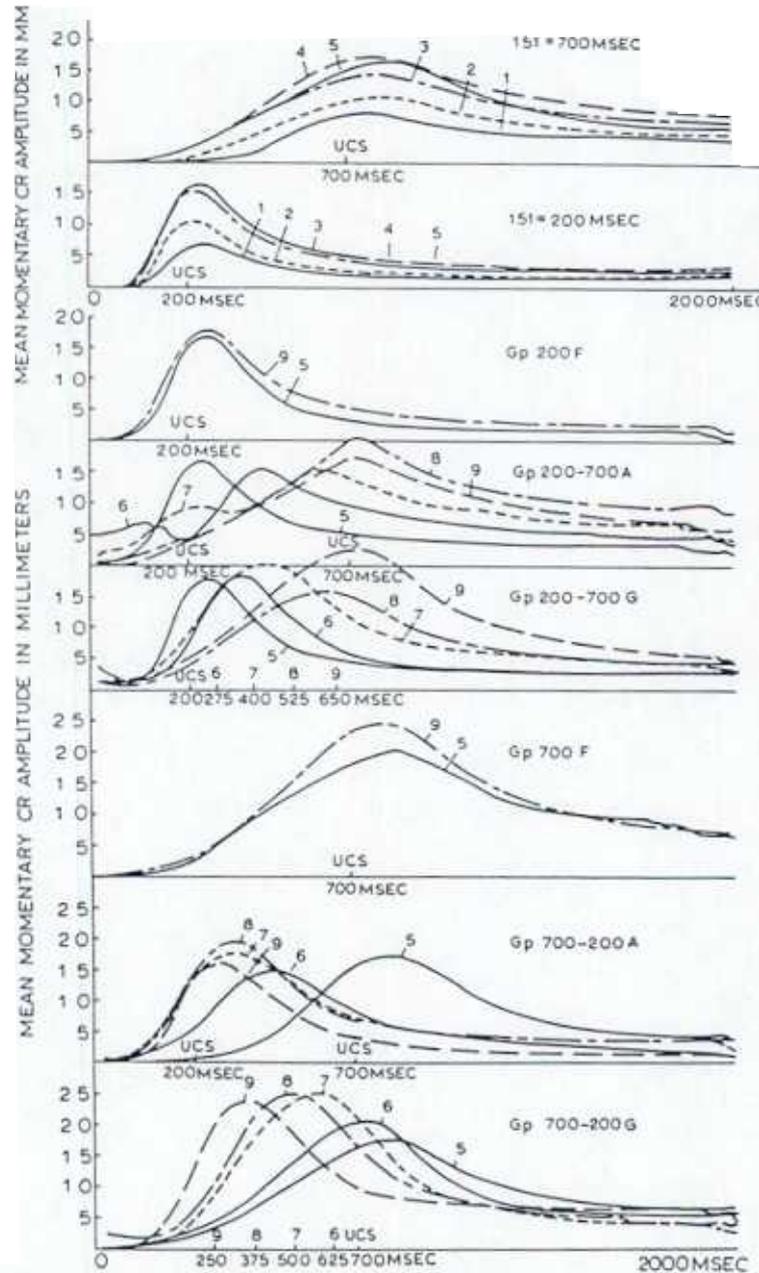


Fig. 20. In the Coleman and Gormezano⁴³ experiments, rabbits were initially conditioned to respond to a 90-dB, 1,000-Hz tone CS by pairing it to a 4-mA, 50-ms 60-Hz periorbital shock US. 72 young (80–100 day old) rabbits were divided into 6 groups. Members of 3 of these groups were initially conditioned to respond with an ISI of 200 ms while members of the other 3 groups were conditioned to respond to stimuli with an ISI of 700 ms. After 5 days of initial conditioning, during which time all of the animals acquired a strong response to the CS beginning shortly after its onset and peaking at roughly the time of onset of the US, the ISIs of some of the subjects were changed. Two groups were exposed to a sudden change from one of the ISIs to the other (200 → 700A and 700 → 200A). Two of the groups were exposed to a gradual shift from one of the ISIs to the other, again symmetrically (200 → 700G and 700 → 200G). As controls, the ISIs of two groups (200C and 700C) were held constant. Subjects' responses were recorded over a period of 4 days in the experimental condition. The nictitating membrane responses under these ISI shift conditions are displayed. (Reprinted from Coleman and Gormezano⁴³ with permission.)

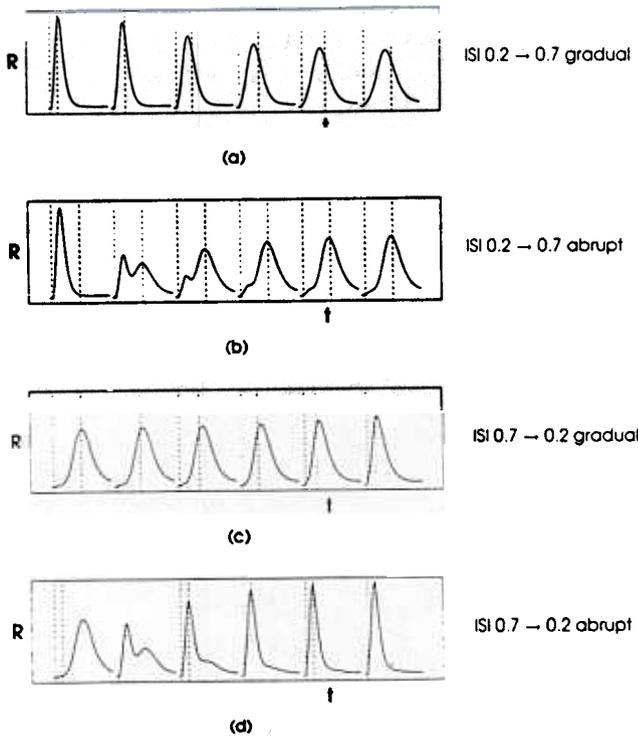


Fig. 21. Simulation of the model output $R(t)$ under ISI shift conditions similar to those used in Coleman and Gormezano⁴³. In the first column of each panel, the model output $R(t)$ is displayed after 25 learning trials. Successive columns are displayed after a block of 8 more learning trials. The vertical lines denote the ISI on the corresponding trial. a: gradual increase of ISI from 0.2 to 0.7 time units on successive learning trials. b: abrupt increase of ISI from 0.2 to 0.7 time units. c: gradual decrease of ISI from 0.7 to 0.2 time units. d: abrupt decrease of ISI from 0.7 to 0.2 units.

level, the hippocampal trace starts forming when responses start being generated and continues to grow stronger through the first 100–150 trials⁹. These data are consistent with the hypothesis that at least part of the hippocampal trace is endogenously generated. It needs also to be noted that, although two of the most important projections to the hippocampus arise from the medial septum and the entorhinal cortex, other

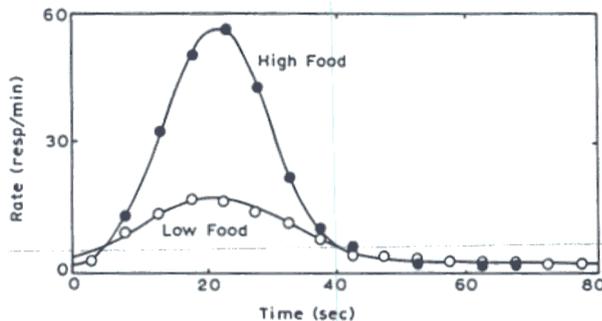


Fig. 22. Effects of differing probability of reinforcement upon the timing and level of response. There is no statistically significant difference between the time at which responding peaks. (Reprinted from Roberts¹⁴⁴ with permission.)

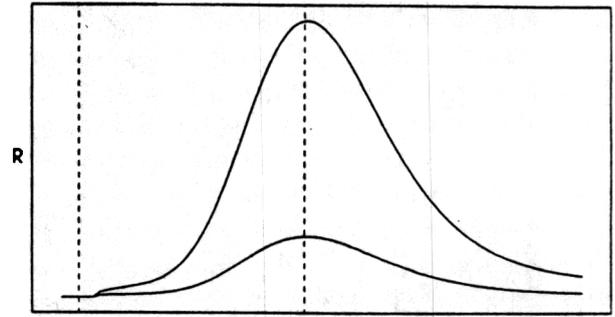


Fig. 23. Model output in a condition simulating the partial reinforcement paradigm of Roberts¹⁴⁴. In the upper curve, 80% of all presentations of the CS were followed by presentations of the US. In the lower curve, 20% of all presentations of the CS were followed by presentations of the US. Despite this difference, the time at which the peak outputs occur is roughly equal, and only the relative levels of output are different.

areas do send projections there too, among them the supermammillary region and the dorsal diagonal gyrus, which project to region CA3 of the hippocampus via the fornix, and the anterior and medial dorsal thalamic nuclei, which project to region CA1 of the hippocampus via the cingulum. Collingridge and Davies (ref. 44, p. 130) discussed additional evidence that “an increase in transmitter release maintains (hippocampal) LTP. This evidence is strongest for the perforant path input

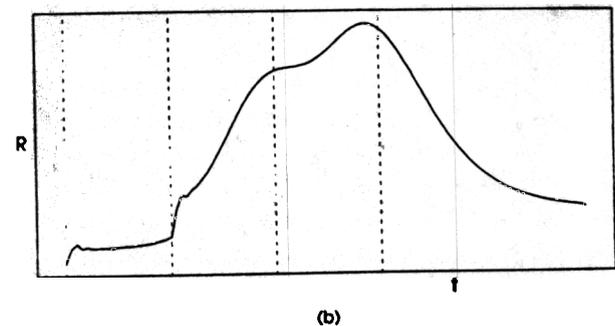
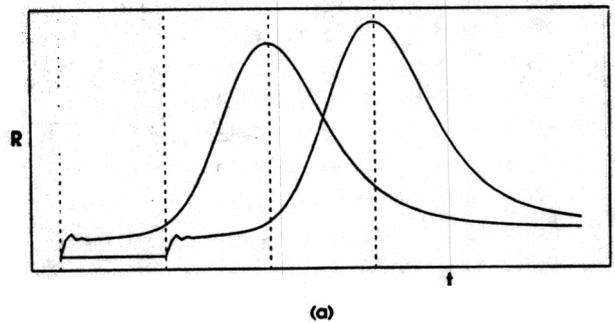


Fig. 24. Simulated time-averaging behavior in response to two CSs. a: model outputs $R(t)$ in response to the two CSs presented individually. b: composite model output $R(t)$ in response to the two CSs presented on the same trail at the same relative onset times as in a. The first vertical line designates the onset time of the first CS, and the second that of the second CS.

to the dentate gyrus, where the increase in transmitter release is dependent on the activation of NMDA receptors.”

20. CONDITIONING AT DENDRITIC SPINES OF DENTATE GRANULE CELLS

The formal processing stages of the START model have a natural hippocampal interpretation that mimics the observed differences between dentate granule cells and hippocampal pyramidal cells during conditioning, and uses a learning mechanism at the model analog of granule cells that is interpreted below in terms of a learned control of transmitter release, with associated alterations in protein synthesis.

The combinatorial explosion of cells that was described in Section 9 is avoided by assuming that the spectral activations x_{ij} are local potentials at the dendritic spines of hippocampal dentate granule cells. Thus the x_{ij} do not correspond to separate cells, but rather to dendritic spines of a single cell that accumulates signals from many sensory representations. The many *pathways* from different sensory representations to the dendrites still need to exist, but their targets are a much smaller population of cells and their dendrites. The microscopic biophysical details of this interpretation will be developed elsewhere. Here we show how the formal linkage of spectral learning properties to hippocampal circuitry leads to new explanations and predictions about hippocampal anatomy and neurophysiology.

In this interpretation, there exists a subset of dentate granule cells that reacts at a single spectral averaging rate r_j in (5), and different subsets of granule cells react at different rates r_j . Each such cell possesses a large number of dendrites that are densely encrusted with dendritic spines. Each spine is assumed to structurally realize a private channel at which individuated activations x_{ij} can be processed at the rate r_j . Learning is activated by a Now Print signal N that globally activates the entire cell. Fig. 25b indicates that the twice-gated operation $f(x_{ij})y_{ij}z_{ij}$ in (11) may be realized in several different ways.

21. CONVERGENCE OF DENTATE GRANULE CELLS AT CA3 PYRAMIDAL CELLS

This interpretation of the START model suggests that (1) conditioning occurs at dentate granule cells, (2) the latency of conditioned firing is constant at individual granule cells, and (3) the hippocampal pyramidal cells to which dentate cells project form a ‘temporal

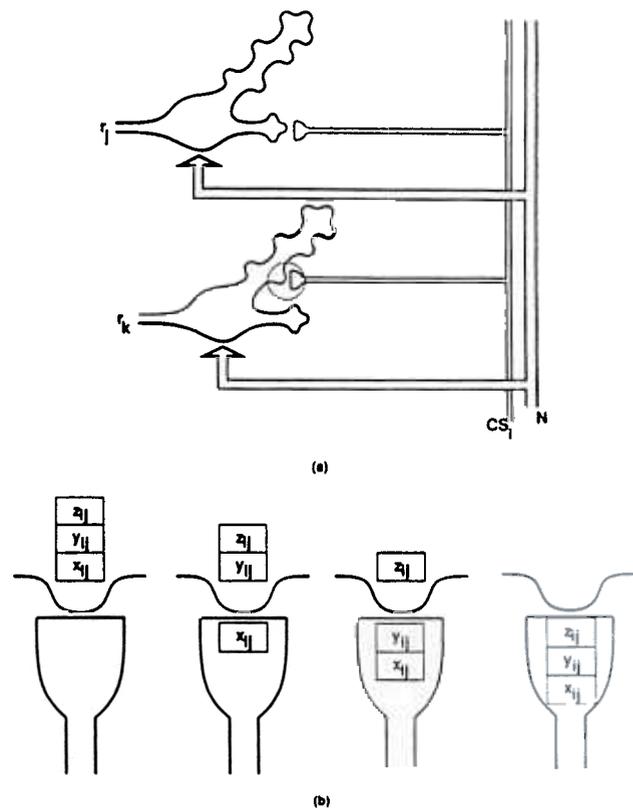


Fig. 25. A possible synaptic realization of spectral timing operations. a: different rates r_j can be realized at different (populations of) dentate granule cells. Each CS_i activates branching pathways whose collaterals synapse at dendritic spines across a subset of cells that include all the rates r_j . The Now Print signal N is delivered in a way that can influence all active synapses across all the dendritic spines. The successive stages x_{ij} , y_{ij} and z_{ij} of cellular activation and gating can, in principle, occur either postsynaptically, or through a combination of presynaptic and postsynaptic operations, as in b.

model’ of adaptively timed behavioral responses. These data are consistent with the model hypothesis, formalized in equation (11), that the individual terms $f(x_{ij})y_{ij}z_{ij}$, which correspond to each fixed and different rate r_j , summate to generate an adaptively timed model R of the behavioral response. We interpret the cells corresponding to different values of r_j as different (subsets of) dentate granule cells, and the cells corresponding to the output R as CA3 pyramidal cells (Fig. 26). It is also assumed that different subsets of CA3 pyramidal cells correspond to different drive representations¹³².

This interpretation of (11) suggests that many dentate granule cells converge on individual CA3 pyramidal cells. This property is consistent with the fact that, in the rat, there are approximately 1,000,000 dentate granule cells but only 160,000 CA3 pyramidal cells^{20,21,162}. In addition, a CA3 cell receives approximately 80 mossy-fiber inputs from dentate granule cells¹⁶². It may thus not be a coincidence that the

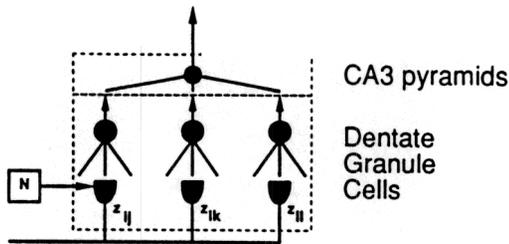


Fig. 26. Interpretation of the output signal $R(t)$ in equation (11) in terms of convergence of dentate granule cell outputs on CA3 pyramidal cells.

Grossberg and Schmajuk⁹⁴ computer simulations and our own found that 80 values of r_j provide an excellent fit to behavioral data on the conditioned NM response. On the other hand, in unreported simulations we have shown that qualitative model properties are robust when the number of populations is increased or decreased by a factor of 4. In any case, the order of magnitude between anatomical and model convergence is acceptable. This anatomical interpretation can be unequivocally tested in terms of the following

Prediction: sets of dentate cells, or perhaps a combination of dentate cells and mossy fibers, exhibit dynamics capable of giving rise to a full spectrum of activation rates r_j .

Gray (ref. 65, pp. 97–100) has surveyed experiments that are consistent with this neurophysiological interpretation of START learning. These data show habituation in dentate granule cells to stimulation of the perforant path, analogous to y_{ij} habituation in response to inputs S_i in equations (5) and (7); potentiation at the dentate synapse in response to perforant path stimulation, analogous to z_{ij} conditioning in response to inputs S_i in equations (5) and (8); swelling in the dendritic spines of dentate granule cells after tetanizing stimulation of the perforant pathway, analogous to the anatomical interpretation of the $f(x_{ij})y_{ij}z_{ij}$ process at dentate spines; and facilitation of dentate response to a perforant path pulse by a prior conditioning pulse to the septum, analogous to the action of the Now Print signal on conditioning in equation (8).

The START model is also consistent with more recent data concerning the effects of manipulations of the dentate gyrus upon the behavior of animals. Diaz-Granados et al.⁵¹ showed that selective dysgenesis of the dentate gyrus in rats due to neonatal X-irradiation impaired performances in a differential-reinforcement-of-low-rate-of-responding (DRL) task: After X-irradiation, which selectively prevents the formation of granule cells, animals were unable to run slowly down a hallway to receive reward. In a similar paradigm, Robinson¹⁴⁶ impaired acquisition of the conditioned

NM response, and Thompson and Disterhoft¹⁶⁶ showed that NMDA agonists and antagonists have opposite effects upon long-interval trace eye blink conditioning. These results are consistent with the proposed interpretation of the START model: animals without functional dentate granule cells, or with impaired NMDA receptors, should be unable to adaptively time their conditioned responses.

This interpretation is also consistent with the lack of effect of granule cell dysgenesis³ or NMDA antagonist treatment¹⁶⁵ upon place learning. The START model suggests that the reinforcing value of an event is less affected by these manipulations than is the ability to adaptively time reinforced behavior. Given this interpretation, the close temporal relationship between being in a place and getting rewarded there may be spared, but distant temporal relationships may be unbridgeable. This possibility may be further testable in the context of fear conditioning, where NMDA antagonists impair fear conditioning over an ISI of 4 s⁴⁸. If this failure is substantially due to a failure of adaptive timing, then near-normal fear conditioning may be found over sufficiently short ISIs if NMDA antagonists are selectively applied to dentate granule cells.

The occurrence of associative learning on dendritic spines also helps to explain how the read-out (or performance) of old associative memories can be decoupled from the read-in (or learning) of new associative memories. Such a dissociation is needed to solve the self-printing problem (Section 5). By (11) read-out of old associative memories is accomplished by the twice-gated signals $f(x_{ij})y_{ij}z_{ij}$; also see Fig. 26. These signals need to be separated from the influence of twice-gated signals activated by other conditioned stimuli and other spectral averaging rates. Dendritic spines can provide this functional separation during read-out, while also being responsive to more global events, such as the Now Print signal N , during LTM read-in via equation (8).

The hypothesis of the START model that hippocampal LTP occurs at dendritic spines in order to functionally dissociate the read-out of old associative memories from the read-in of new associative memories was discussed in Grossberg⁷⁶ (see Fig. 25). This type of process has recently excited a great deal of further work based upon new experimental approaches to hippocampal LTP and the discovery of the NMDA receptor^{25,26,53}. The START model is consistent with data showing that conditioning takes place at NMDA receptors in the perforant-to-dentate pathway. As Collingridge and Davies (ref. 44, p. 130) have noted: "Most neurochemical evidence suggests that an increase in transmitter release maintains LTP. This evi-

dence is strongest for the perforant path input to the dentate gyrus, where the increase in transmitter release is dependent on the activation of NMDA receptors.”

22. NMDA RECEPTORS AND ADAPTIVE TIMING

The recent experiments suggesting that an increase in presynaptic transmitter release may help to control LTP at dentate granule cells and includes activation of NMDA receptors^{44,55} are consistent with another early prediction^{67,68} about associative learning. This prediction suggested that associative learning is achieved by “joint control of presynaptic excitatory transmitter production by presynaptic and postsynaptic levels of membrane potential. This control is presumed to be effected by the interaction of the pairs (Na^+ , K^+) and (Ca^{2+} , Mg^{2+}) of antagonistic ions whose binding properties to intracellular sites and enzymes set cellular production levels” (ref. 68, p. 325). In particular, a synergetic interaction of a voltage-dependent, postsynaptically generated, inward Ca^{2+} current with inward Na^+ and outward K^+ currents was predicted, as well as a competitive interaction between Ca^{2+} and Mg^{2+} . Recent studies of LTP at NMDA receptors have reported and greatly elaborated contemporary understanding of this sort of interaction, including the competition between Ca^{2+} and Mg^{2+} (ref. 44).

Related predictions may now be testable at the perforant path—CA3 pyramidal cell synapse. One prediction suggests certain “nerve cells are capable of learning as ‘chemical dipoles’” (ref. 68, p. 325) that control the availability of the proper relative amounts of Ca^{2+} , Mg^{2+} , Na^+ , and K^+ , among other chemicals, at the cell sites where they are needed. Such control is suggested to coordinate potentiation of presynaptic transmitter production with levels of postsynaptic protein synthesis aimed at enabling the postsynaptic cell to cope with time-varying loads of presynaptic input. Akers et al.¹ have shown that “protein kinase activation leading to phosphorylation of neural proteins appears to occupy a pivotal role in the development and expression of synaptic plasticity” in response to perforant path stimulation (p. 587). Further experiments are needed to test possible correlations between presynaptic and postsynaptic effects. This proposal also suggested that the shape of neurons realizes a type of structural dipole that helps to support the dynamics of the chemical dipole. The two poles of the structural dipole, at the dendritic apparatus/cell body and the synaptic knobs, respectively, were suggested to help maintain chemical gradients along the axons between

these poles. Interactions between the cell nucleus, the cell membrane, microtubules and tight junctions between presynaptic terminals and postsynaptic cells were proposed to maintain these gradients. Further details concerning these predicted chemical dipole properties may be found in Grossberg⁶⁸.

PART II. REINFORCEMENT, RECOGNITION AND MOTOR LEARNING

The spectral timing part of the circuit in Fig. 5 is new. The remainder of the circuit is part of a larger theory concerning the neural substrates of reinforcement, recognition, attention, memory search and motor control. Relevant parts of the theory are summarized below. They are used to clarify how the adaptive timing circuit interacts with other types of brain circuits, and to show how recent neurobiological data support the existence of each of the model’s processing stages.

23. REINFORCEMENT LEARNING IN VERTEBRATES AND INVERTEBRATES

In Section 4, a drive representation D was defined as a population of cells at which sensory, reinforcement and homeostatic, or drive, signals converge to regulate reinforcement learning, emotional reactions and motivational decisions. Fig. 27 depicts the type of model circuit in which drive representations were described in

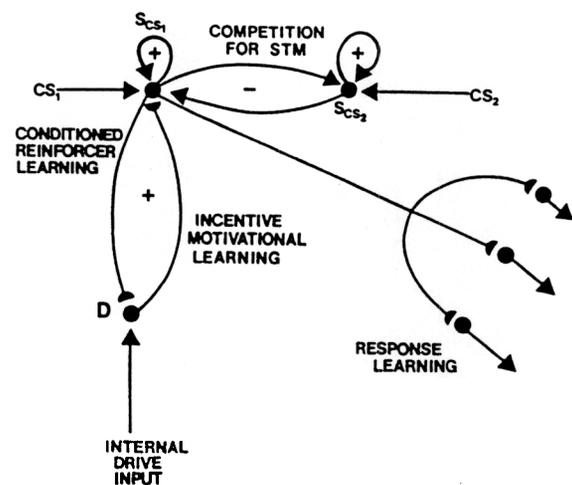


Fig. 27. Schematic conditioning circuit: conditioned stimuli (CS_i) activate sensory representations (s_{CS_i}) which compete among themselves for limited capacity STM activation and storage. The activated s_{CS_i} elicit conditionable signals to drive representations and motor command representations. Learning from a 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 external sensory plus internal drive inputs is sufficiently large.

Grossberg⁷¹ to explain vertebrate conditioning data. A similar model (Fig. 28) has recently been used to explain invertebrate conditioning data from experiments on *Aplysia*^{28,84}. The use of a similar circuit by such different species is clarified by the fact that it is the simplest solution of two general learning problems, called the synchronization problem and the persistence problem, that all animals capable of classical conditioning in an unconstrained environment need to solve^{71,82,85}. The synchronization problem asks how classical conditioning can occur without massive interference if the ISI on each learning trial may be different and irrelevant stimuli may occur between the CS and US. The persistence problem asks how alternating CSs, each conditioned to different emotional responses, are protected against rapid extinction due to association with the 'wrong' emotional response.

Fig. 27 contains pathways that were omitted from Fig. 5 for simplicity. As noted in Section 4, during classical conditioning, pairing of a CS₁ sensory representation s_{CS_1} with activation of a drive representation D by a reinforcer US causes the modifiable synapses connecting s_{CS_1} with D to become strengthened. This conditioning event converts CS₁ into a conditioned reinforcer. Fig. 27 shows reciprocal conditionable pathways from the drive representations D to the sensory representations s . Conditioning of these pathways is called incentive motivational learning. Activation of conditioned $s \rightarrow D \rightarrow s$ feedback pathways by CS₁ can shift attention towards the set of all previously reinforced sensory cues that are motivationally consistent with D .

This shift of attention occurs because the sensory representations, in addition to emitting conditioned reinforcer signals and receiving incentive motivation signals, compete among themselves (Fig. 27) for a limited capacity STM. When strong incentive motivational feedback signals are received at the sensory

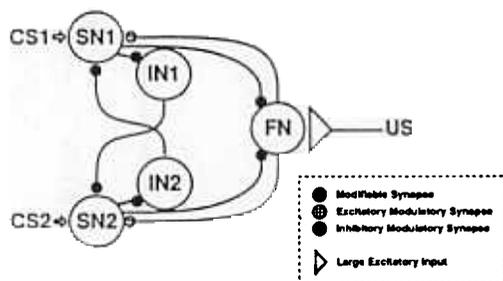


Fig. 28. A model for conditioning in *Aplysia*. SN, sensory neuron; FN, facilitatory neuron; IN, inhibitory neuron. The SNs play the role of sensory representations s , the FNs the role of a drive representation D , and the INs carry out the competition between sensory representations.

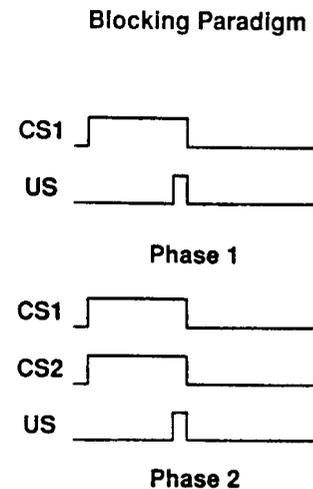


Fig. 29. A schematic of the Pavlovian blocking paradigm. The two phases of the experiment are discussed in the text.

representational field, these signals can bias the competition for STM activity towards the set of motivationally preferred cues.

24. ATTENTIONAL FOCUSING, BLOCKING AND THE HIPPOCAMPUS

The feedback signals $s \rightarrow D \rightarrow s$ generate a resonant state of activation between levels s and D that focuses attention upon recognition codes of events which have led to reinforcing consequences during past experiences. Such attentional modulation enables a biological information processing system to selectively process those environmental inputs that are most important to its current goals. A typical example of such selective processing is illustrated by the blocking paradigm shown in Fig. 29¹¹⁰. First, a conditioned stimulus CS₁, such as a tone, is presented several times, followed at a given time interval by an unconditioned stimulus US, such as electric shock, until a conditioned response CR, such as fear, develops. Then CS₁ and another conditioned stimulus, CS₂, such as light, are presented simultaneously, followed at the same time interval by the US. After conditioning, CS₂ is presented alone, not followed by a US, and no CR occurs. Intuitively, CS₁ 'blocks' conditioning of the simultaneously presented CS₂ because CS₁, by itself, perfectly predicts its consequence, the US. The CS₂ is thus redundant and unpre- dictive, hence does not get conditioned to the US.

The blocking property may be explained in terms of 4 properties of the network in Fig. 27: (1) pairing of a CS₁ with a US in the first phase of the blocking experiment endows the CS₁ cue with properties of a conditioned, or secondary, reinforcer; that is, the positive feedback pathway $s_1 \rightarrow D \rightarrow s_1$ between the drive

representation D and the sensory representation s_1 of CS_1 is strengthened due to learning. (2) These reinforcing properties of a CS_1 shift the focus of attention towards its own processing at s_1 . (3) The processing capacity of attentional resources is limited, as a result of the competition between sensory representations s . Thus a shift of attention towards one set of stimuli can prevent other stimuli, such as CS_2 , from being attended (Fig. 27). Withdrawal of attention from the sensory representation s_2 of the stimulus CS_2 prevents that representation from entering new conditioned relationships, by attenuating learning from s to D and from D to s . Learning is attenuated when the activity S_2 of s_2 becomes small, because it is regulated by an activity-dependent gate, as in (4).

Just as simultaneous occurrence of a conditioned reinforcer CS_1 with a new CS_2 can block conditioning of CS_2 , so too can simultaneous occurrence of a primary reinforcer US with a new CS block conditioning of CS. This latter property helps to explain why US onset needs to occur after CS onset in order for effective conditioning to occur^{28,82,91}.

One way to verify whether a neural model has processing stages that correlate well with brain circuits is to test if a formal model lesion has effects similar to those of a corresponding brain lesion on behavioral properties. Grossberg (ref. 76, Fig. 24) suggested that a final common path within (an expanded model of) a drive representation D includes the hippocampal formation. Eliminating the 'hippocampal formation' in the model would therefore weaken $D \rightarrow s$ feedback signals, and thus the model's mechanism of blocking. Hippocampal lesions do, in fact, prevent blocking from occurring. Both CS_1 and CS_2 can be conditioned in a blocking experiment performed on a hippocampectomized animal^{116,143,150,156}. Likewise, hippocampectomized animals find it hard to actively ignore non-reinforced cues¹⁴².

These experiments also showed that hippocampal lesions do not interfere with emotional conditioning. Although such a dissociation could not be explained in the model of Figs. 27 and 28, it can be explained using the model of Fig. 5, which distinguishes the circuit for adaptive timing from the circuit for emotional conditioning. A circuit which combines the components of Fig. 5 with those of Fig. 27 is shown in Fig. 30.

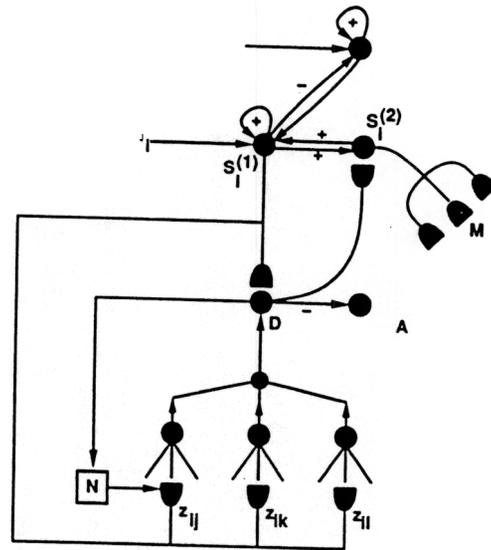


Fig. 30. A START model with feedback pathways $D \rightarrow s^{(2)} \rightarrow s^{(1)}$ that are capable of focussing attention in an adaptively timed fashion on reinforcing events. The sensory representations s are broken into two successive levels $s^{(1)}$ and $s^{(2)}$. Levels $s^{(1)}$ and $s^{(2)}$ interact via reciprocal excitatory pathways. The excitatory pathways $s^{(1)} \rightarrow D$ and $D \rightarrow s^{(2)}$ are adaptive. Representations in $s^{(2)}$ can 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. Thus, prior to conditioning, a CS can 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 positive $s^{(1)} \rightarrow D \rightarrow s^{(2)} \rightarrow s^{(1)}$ feedback and draw attention to itself as it activates the emotional representations and motivational pathways controlled by D .

was shown that each sensory representation s need to be broken into two successive stages $s^{(1)}$ and $s^{(2)}$, as in Fig. 30, such that $s^{(1)}$ projects to both $s^{(2)}$ and D , D projects to $s^{(2)}$, and $s^{(2)}$ projects back to $s^{(1)}$. The pathway $s^{(1)} \rightarrow D$ can support emotional conditioning, thereby converting the stimuli represented at $s^{(1)}$ into conditioned reinforcers. The pathway $D \rightarrow s^{(2)}$ supports incentive motivational conditioning. It primes all sensory representations that are motivationally consistent with D . The multisynaptic pathway $D \rightarrow s^{(2)} \rightarrow s^{(1)}$ provides the feedback from D that supports the blocking process. This expanded version of the model has been used to computationally simulate blocking data⁹¹.

Grossberg^{76,79} interpreted the first stage of sensory processing $s^{(1)}$ as a thalamic representation, the second stage $s^{(2)}$ as a cortical representation, the first stages of drive representational processing D as networks, such as hypothalamus and amygdala, that are involved in homeostatic and emotional processing, and the final stages of drive representational processing as including the hippocampal formation. With this interpretation, the conditioning of $s^{(1)} \rightarrow D$ synapses in Fig. 30 predicts that subcortical emotional conditioning is possible. This prediction has been supported by recent experiments

25. SUBCORTICAL FEAR CONDITIONING AND A CORTICAL ROLE IN EXTINCTION

The circuit in Fig. 30 includes sensory representations s that process incoming signals in two successive processing stages $s^{(1)}$ and $s^{(2)}$. In Grossberg^{71,76,79,82} it

They also summarized data from their laboratory showing that “the dorsal accessory olive-climbing fiber projection is the necessary and sufficient US pathway... the mossy fiber projection is the necessary and sufficient CS pathway... and... appropriately timed conjoint activation of mossy fibers as the CS and climbing fibers as the US yields normal learning of discrete, adaptive behavioral CSs” (ref. 167, pp. 387–388). Clark et al.⁴¹ have shown, moreover, that lesions of the deep cerebellar nuclei not only abolish the conditioned behavioral response, but also abolish the hippocampal temporal model of the behavior (Part I) in response to ipsilateral USs. Administration of contralateral USs quickly causes the reacquisition of both the neuronal model and the behavioral response. Taken together, these data strongly support the hypothesis that conditioning of the CS-activated discrete adaptive response occurs in the cerebellum and that these cerebellar signals are needed for expression of the hippocampal temporal model.

The hypothesis that the cerebellum helps to control motor learning has a long history. Brindley²⁴ and Grossberg⁶⁶ were among the first to model motor learning in the cerebellum at the synapses between cerebellar parallel fibers and Purkinje cell dendritic spines, using the climbing fibers as a teaching signal. Grossberg⁷⁰, Marr¹²⁴, and Albus² further modelled this concept. Marr¹²⁴ suggested that these synapses increase in strength due to learning; Albus² suggested that they decrease in strength; Grossberg⁷⁰ suggested that they may either increase or decrease in strength, depending upon the learning context. Subsequent models of cerebellar motor learning include those of Bullock and Grossberg²⁷, Fujita^{57,58}, Grossberg and Kuperstein^{89,90}, Houk et al.¹⁰² and Ito^{104,105}. These cerebellar models have been used to analyze behavioral and neural data about eye and arm movements, such as the results of Ebner and Bloedel⁵², Gilbert and Thach⁶², Ito¹⁰⁵, Optican and Robinson¹³³ and Ron and Robinson¹⁴⁷. In addition to their discussions of nictitating membrane and jaw movement conditioning, Thompson et al.^{167,168} also summarized experiments demonstrating motor learning in the cerebellum during classical conditioning of the limb flexion reflex. Thus the cerebellum plays a key role in conditioning motor responses of eye, arm, leg, nictitating membrane and jaw movements, among others.

These recent data and models about cerebellar learning clarify how motor responses are adaptively controlled, and also suggest that motor learning differs from the types of conditioned reinforcer learning, incentive motivational learning, recognition learning and adaptive timing that are depicted in Figs. 5, 27 and 30.

27. MACROCIRCUIT FOR SENSORY-COGNITIVE PROCESSING: ADAPTIVE RESONANCE THEORY

It remains to describe how the orienting system A in Fig. 5 is controlled; in particular, how the unexpected non-occurrence of a reinforcer can activate A and thereby cause orienting reactions, attention shifts and emotional frustration. With this information in hand, the hypothesis that drive representations D inhibit A in response to expected non-occurrences can be better understood as a mechanism for preventing maladaptive reactions to predictive cues. It also remains to discuss how extinction is controlled via this process, as remarked in Section 25. For this, see refs. 82, 83 and 93.

These types of reactions are modelled by sensory-cognitive circuits that are called Adaptive Resonance Theory, or ART models. ART models have been used to explain and predict a large body of cognitive and neural data about recognition learning, attention and memory search^{36,45,81,85–87}. ART systems suggest a solution to a fundamental learning problem that is called the *stability-plasticity dilemma*: an adequate self-organizing recognition system must be capable of *plasticity* in order to learn about significant new events, yet it must also remain *stable* in response to irrelevant or often repeated events. In order to prevent the relentless degradation of its learned codes by the ‘blooming, buzzing confusion’ of irrelevant experience, an ART system is sensitive to *novelty*. It is capable of distinguishing between familiar and unfamiliar events, as well as between expected and unexpected events.

The importance of expectancy-related processes in conditioning and cognitive processes has been extensively documented since the pioneering work of Tolman¹⁶⁹ and Sokolov^{154,155}. In ART, interactions between an attentional subsystem and an orienting subsystem, or novelty detector, enable the network to self-stabilize its learning, without an external teacher, as the learned recognition code becomes globally self-consistent and predictively accurate; in other words, as the system familiarizes itself with an environment by categorizing the information within it in a way that leads to behavioral success. The attentional subsystem undergoes both bottom-up learning and top-down learning within the LTM-marked pathways between the processing levels denoted by F_1 and F_2 in Fig. 31. The top-down LTM process learns expectations. The network self-stabilizes its learning by matching its top-down expectations against bottom-up input patterns and using the degree of match or mismatch to regulate processes of learning or memory search, respectively. Thus ART suggests how novelty-sensitive matching processes regulate the course of learning.

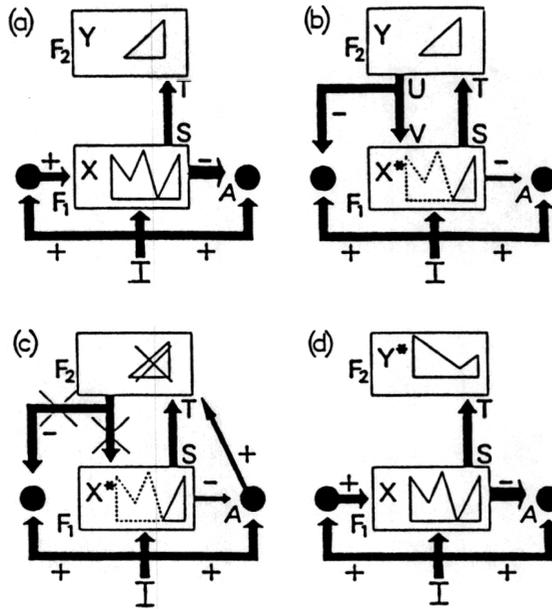


Fig. 32. ART search for an F_2 code. a: the input pattern I generates the specific STM activity pattern X at F_1 as it non-specifically activates the orienting subsystem A . Pattern X both inhibits A and generates the output signal pattern S . Signal pattern S is transformed into the input pattern T , which activates the STM pattern Y across F_2 . b: Pattern Y generates the top-down signal pattern U which is transformed into the prototype pattern V . If V mismatches I at F_1 , then a new STM activity pattern X^* is generated at F_1 . The reduction in total STM activity which occurs when X is transformed into X^* causes a decrease in the total inhibition from F_1 to A . c: if the matching criterion fails to be met, A releases a non-specific arousal wave to F_2 , which resets the STM pattern Y at F_2 . d: after Y is inhibited, its top-down prototype signal is eliminated, and X can be reinstated at F_1 . Enduring traces of the prior reset lead X to activate a different STM pattern Y^* at F_2 . If the top-down prototype due to Y^* also mismatches I at F_1 , then the search for an appropriate F_2 code continues.

By itself, the attentional subsystem is unable simultaneously to maintain stable representations of familiar categories and to learn new categories for unfamiliar patterns. An isolated attentional subsystem may be either rigid and incapable of creating new recognition categories for unfamiliar patterns, or unstable and capable of ceaselessly recoding the recognition categories of familiar patterns.

The orienting subsystem interacts with the attentional subsystem to maintain the stability-plasticity balance. It resets the STM of the attentional subsystem when an unfamiliar event occurs and thereby initiates a memory search within the attentional subsystem for a better internal representation. The orienting subsystem is thus essential for expressing whether an input pattern is familiar and well represented by an existing recognition code, or unfamiliar and in need of a new recognition code.

Fig. 32 illustrates a typical ART memory search cycle. As shown in Fig. 32a, an input vector I registers

itself as a pattern X of activity, or short term memory (STM), across level F_1 . The F_1 output vector S is then transmitted through the multiple converging and diverging adaptive pathways emanating from F_1 . This transmission event multiplies the vector S by a matrix of adaptive weights, or long term memory (LTM) traces, to generate a net input vector T to level F_2 . Lateral inhibitory interactions within F_2 contrast-enhance vector T . A compressed activity vector Y is thereby generated across F_2 .

Activation of F_2 nodes may be interpreted as 'making a hypothesis' about an input I . When Y is activated, it generates a signal vector U that is transmitted along top-down adaptive pathways. After multiplication of these top-down signals by a matrix of adaptive weights, or LTM traces, a net vector V inputs to F_1 (Fig. 32b). Vector V plays the role of a learned top-down expectation. Activation of V by Y may be interpreted as 'testing the hypothesis' Y , or 'reading out the category prototype' V . ART networks are designed to match the 'expected prototype' V of the category against the active input pattern, or exemplar, I .

This matching process may change the F_1 activity pattern X by suppressing activation of all the feature detectors in I that are not confirmed by V . The resultant pattern X^* encodes the pattern of features to which the network 'pays attention'. If the expectation V is close enough to the input I , then a state of *resonance* occurs as the attentional focus takes hold. Damasio⁴⁷ uses the term 'convergence zones' to describe the process whereby an activation pattern X^* across distributed features is bound together by resonant feedback. The resonant state persists long enough for learning to occur; hence the term *adaptive resonance* theory. ART systems learn prototypes, rather than exemplars, because the attended feature vector X^* , rather than the input I itself, is learned.

The criterion of an acceptable match is defined by a dimensionless parameter called *vigilance*. The vigilance parameter is computed in the orienting subsystem A , where it may be increased by punishing events or other unexpected consequences^{31,38,39}. Vigilance weighs how close the input exemplar I must be to the top-down prototype V in order for resonance to occur. Because vigilance can vary across learning trials, recognition categories capable of encoding widely differing degrees of generalization, or morphological variability, 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**. In the limit of very high vigilance, prototype learning reduces to exemplar learning. Thus a single ART system may be

used, say, to recognize abstract categories of faces and dogs, as well as individual faces and dogs.

If the top-down expectation V and the bottom-up input I are too novel, or unexpected, to satisfy the vigilance criterion, then a bout of hypothesis testing, or memory search, is triggered. Memory search leads to selection of a better recognition code at level F_2 with which to represent input I at level F_2 . The *orienting subsystem* A mediates the search process. During search, the orienting subsystem interacts with the attentional subsystem, as in Fig. 32c,d, to enable the attentional subsystem to learn about novel inputs without risking unselective forgetting of its previous knowledge.

The search process prevents associations from forming between Y and X^* if X^* is too different from I to satisfy the vigilance criterion. The search process resets Y before such an association can form. A familiar category may be selected by the search if its prototype is similar enough to the input I to satisfy the vigilance criterion. The prototype may then be refined in light of new information carried by I . If I is too different from any of the previously learned prototypes, then an uncommitted population of F_2 cells is selected and learning of a new category is initiated.

A network parameter controls how deeply the search proceeds before an uncommitted node is chosen. As learning of a particular category self-stabilizes, all inputs coded by that category access it directly, without the need for search. Familiar, consolidated memories can thus be accessed in a one-pass fashion, after resetting the previously active category. The category selected is the one whose prototype provides the globally best match to the input pattern. In a situation where a mixture of familiar and unfamiliar events are experienced, familiar inputs can directly activate their learned categories, while novel inputs continue to trigger adaptive memory searches for better categories, until the network's memory capacity is fully utilized.

These ART mechanisms include the processes that are needed to interpret the effects of $D \rightarrow A$ inhibition that were described in Section 1. These include a process whereby learned expectations may be mismatched by a sensory expectation at level F_1 of the attentional subsystem in Fig. 31. When a mismatch of bottom-up exemplar and top-down prototype occurs, the orienting subsystem is activated, giving rise to a STM reset wave in the form of a non-specific arousal burst to the attentional subsystem (Fig. 32c). This arousal burst acts to reset the sensory representations of all sensory events that are currently active in STM within the attentional subsystem. Representations with high STM activation tend to become less active, repre-

sentations with low STM activation tend to become more active, and the novel event which caused the mismatch tends to be more actively stored than it would have been had it been expected. Banquet and Grossberg⁷ have discussed experiments on human event-related potentials (ERPs) during probabilistic choice reaction time tasks that have tested the predicted chronometry of the mismatch-arousal-reset sequence in terms of the P120-N200-P300 sequence of ERPs. One effect of STM reset is to shift the focus of attention towards sensory representations which may better predict environmental contingencies. In a classical conditioning paradigm, such an attention shift can dishabituate, or unblock, sensory representations that were not attended before the STM reset event^{81,82}. Activation of the orienting subsystem also triggers orienting responses, such as the activation of motor reactions to orient towards the unexpected event (Fig. 31).

This organization of learned expectations, attention shift mechanisms, and orienting mechanisms within ART allowed Grossberg and Schmajuk⁹⁴ to hypothesize that activation of the drive representation D gates, or inhibits, the orienting subsystem A . Activation of this inhibitory gate prevents reset of the attentional focus and release of orienting behaviors if an expected non-occurrence is experienced. Such a gating operation does not, however, prevent a sensory match from being detected earlier than usual, because matches with learned expectations occur within the attentional subsystem, not the orienting subsystem. At times when the adaptive timing mechanism is inactive, the gate is open. Then activation of the orienting subsystem can trigger reset of STM and orienting reactions in response to unexpected non-occurrences.

28. HIPPOCAMPAL LESIONS AND MEDIAL TEMPORAL AMNESIA

The division of labor within ART, between an attentional subsystem and an orienting subsystem, thus provides the type of processing substrate that is needed to instantiate adaptive timing heuristics. This division of labor has also been helpful in clarifying many other types of data. For example, Carpenter and Grossberg^{31,34} have pointed out that a lesion of the ART orienting subsystem creates a memory disturbance whose formal symptoms are similar to those of humans afflicted with medial temporal amnesia, including 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^{42,64,122,125,160,161,173,174,178}.

Unlimited anterograde amnesia occurs in the model because, without a functional orienting subsystem, the network cannot carry out the memory search and subsequent learning needed to establish a new recognition code. Limited retrograde amnesia occurs because familiar events can directly access their recognition codes, without activating the orienting subsystem. Before events become familiar, a period of memory consolidation occurs during which the orienting subsystem does play a role, as indicated in Fig. 32c. This failure of consolidation does not prevent learning per se. Instead, learning is associated with the first recognition category that is activated by bottom-up processing, much as "amnesics are particularly strongly wedded to the first response they learn" (ref. 65, p. 253). Abnormal reactions to novelty, including perseverative reactions, occur. In an ART circuit, this happens because the orienting subsystem cannot carry out its normal function of STM reset, and therefore cannot inhibit sensory representations or top-down expectations that may be persistently mismatched by bottom-up sensory signals. The inability to search memory via its orienting subsystem prevents an ART system from discovering more appropriate stimulus combinations to which to attend. In a similar vein, Butters and Cermak (ref. 29, p. 393) reported that "Korsakoff patients' encoding deficits may be related to a general impairment in their ability to attend to relevant dimensions of stimuli." Normal priming is possible in an ART model because

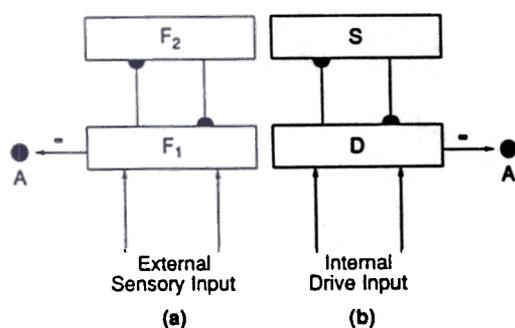


Fig. 33. A schematic representation showing the close homology between the sensory-cognitive ART circuit shown in Fig. 31 and the cognitive-emotional circuit shown in Fig. 30. a: the sensory-cognitive circuit consists of a level F_1 for representing activation of sensory features. Level F_1 interacts with a sensory representation level F_2 that encodes learned chunks, or compressed representations, of the sensory features. Level F_1 interacts with level F_2 via reciprocal pathways that are adaptive and excitatory. Level F_1 also inhibits the orienting subsystem A . b: the sensory-drive ART circuit consists of a drive representation level that interacts with a sensory representation level S via reciprocal pathways that are adaptive and excitatory. Level D also inhibits the orienting systems A . The circuits a and b are combined by incorporating level F_2 into level S as described in the text.

it can be mediated entirely by the attentional subsystem, notably the top-down expectations of this subsystem. The close correspondence between the symptoms of medial temporal amnesia and the formal properties of an ART model with defective orienting subsystem is consistent with accumulating evidence for the hypothesis⁷⁶ that the in vivo analog of the ART orienting subsystem intersects, or is closely linked to, the hippocampal formation.

Similar behavioral problems have been identified in hippocampectomized monkeys. Gaffan⁵⁹ noted, for example, that fornix transection "impairs ability to change an established habit...(there is) impaired learning ability when one habit is to be formed in one set of circumstances and a different habit is to be formed in a different set of circumstances that is similar to the first and therefore liable to be confused with it" (p. 94). A similar problem occurs in an ART network with a defective orienting subsystem. Such a defect prevents STM reset, which normally leads to memory search and learning of different representations for the two similar events. Pribram¹⁴² calls 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 structures¹⁷⁸, and clarify why hippocampectomized rats have difficulty orienting to novel cues¹³¹ and why there is a progressive reduction in novelty-related hippocampal potentials as learning proceeds in normal rats^{49,50}. In summary, localization of both orienting subsystem circuits and adaptive timing circuits in, or intimately related to, the hippocampal formation helps to explain a large body of neuropsychological data. Further hippocampal relationships to ART model mechanisms will be discussed below.

29. A SYNTHESIS OF SENSORY-COGNITIVE AND COGNITIVE-REINFORCER CIRCUITS

We are now ready to join together the sensory-cognitive ART network in Fig. 31a with the cognitive-reinforcer and adaptive timing network in Fig. 30. When this is done, a striking formal similarity between the different types of circuits may be discerned. This similarity suggests that cognitive and emotional processes in the brain share many design properties in common^{82,83}, unlike artificial intelligence models of problem solving.

The sensory representations s in Fig. 30 are recognition codes for sensory events. For definiteness, we identify them with the recognition codes at the level F_2

of the ART network in Fig. 31a, as in Fig. 31b. When this is done, Fig. 30 may be redrawn in a way that reveals a striking homology with the ART recognition circuit in Fig. 31. A comparison between Fig. 33a and b illustrates this homology. In Fig. 33a and b, the sensory representation level F_1 and the drive representation level D play an analogous role. In particular, both level F_1 and level D send inhibitory signals to the orienting subsystem A . The inhibitory signals from F_1 prevent A from resetting STM at level F_2 unless a sensory input pattern mismatches a top-down learned expectation at level F_1 . The inhibitory signals from D help to prevent A from resetting level F_2 when a reinforced event is being attended. As noted in Section 24, such an attentive focus develops due to an exchange of positive feedback signals between levels F_2 and D , supported by conditioned $F_2 \rightarrow D \rightarrow F_2$ pathways.

30. INFLUENCES OF HIPPOCAMPECTOMY ON CONDITIONED TIMING

This synthesis of cognitive and emotional networks enables the theory to explain a broad range of data concerning changes in conditioned timing that are due to hippocampectomy. The expanded model clarifies why the hippocampus is not needed for delay conditioning, but is needed for classical conditioning of more complex associations, such as reversal conditioning and trace conditioning. It has been shown that bilateral hippocampectomy severely disrupts the rate of reversal of two-tone discrimination¹¹, reversal of cross-modality discrimination and tone-light discrimination¹⁷⁵. Hippocampectomy does not, however, disrupt initial learning of the discrimination⁹. The deficit in reversal conditioning is consistent with the explanation of perseverative behavior due to disrupted STM reset and memory search that was given in Section 28.

Hippocampectomy has a profound effect on NM response shape during trace conditioning; for example, if a 100-ms CS duration and a 500-ms ISI are used. Then small, short-latency responses occur to the CS, rather than the large, adaptively timed long-latency responses of control animals¹⁵⁹. The removal of the spectral timing process clarifies why the timed responses are eliminated. A detailed study of the model circuit also clarifies why some responses remain. As Fig. 30 illustrates, the interactions between sensory representations s and drive representations D survive ablation of the adaptive timing circuit, so that certain aspects of motivated behavior remain intact. On the other hand, the role of the adaptive timing circuit in prolonging reactions to sensory cues, and in regulating

the duration of motivated attention, are no longer available. This analysis also clarifies why properties of delay conditioning are altered by hippocampectomy^{9,141}.

The model gains additional support from its ability to rationalize this pattern of conditioned behavioral changes due to hippocampectomy. Although these data strongly suggest that the hippocampus plays an important role in the control of timing, they do not imply that other brain regions do not also contribute to the hippocampally observed 'temporal model'. The very fact that hippocampectomy alters conditioned behavioral timing indicates that timing is conditioned at hippocampal sites, as well as at non-hippocampal sites, such as the sites that control the cerebellar conditioned reactions (Section 26).

31. CONCLUDING REMARKS: VARIETIES OF LEARNING FUNCTIONS AND NETWORKS

The theory developed in this article provides a computational framework in which many behavioral and neural data about conditioning can be analyzed. By identifying several problems that a behaving organism needs to solve in order to survive, the theory has been able to distinguish between several functionally distinct learning processes, to model several of their main mechanisms, and to outline a system architecture within which they are combined. These learning processes include adaptive timing, and the way in which it selectively inhibits inappropriate reactions to expected non-occurrences; reinforcement learning, notably emotional conditioning; incentive motivational learning, including the allocation of attention and the energizing of behavioral responses; recognition learning, including the bottom-up learning that initiates selection of recognition categories and the top-down learning of expectations that help to calibrate novelty and to control memory search; and response learning, including the conditioning of discrete defensive reflexes. This sort of integrative theory exhibits features that are still quite novel in computational neuroscience. This is particularly true where the theory links together several conceptual and organizational levels in order to experimentally support its hypotheses. The theory provides behavioral analyses that help to identify functionally distinct brain processes, mathematically precise circuits that model these processes, unifying design principles to tie these circuits together into a system architecture, neural markers in terms of identifiable anatomical and physiological processes, and computer simulations and predictions to test this architecture at multiple levels of

behavioral and neural organization. With such a foundation in hand, every new datum creates a series of implications that may support or confront the theory at multiple points, thereby creating multiple constraints for propelling further theoretical tests, modifications and refinements. Such theories seem necessary if the immense masses of behavioral and neural data already available are ever to achieve a rational explanation.

Acknowledgements. Thanks to Cynthia E. Bradford and Diana Meyers for their valuable assistance in the preparation of the manuscript and illustrations. S.G. was supported in part by the Air Force Office of Scientific Research (AFOSR 90-0175), DARPA (AFOSR 90-0083), the National Science Foundation (NSF IRI-90-24877), and the Office of Naval Research (ONR N00014-91-J-4100). J.W.L.M. was supported in part by the Air Force Office of Scientific Research (AFOSR 90-0128).

REFERENCES

- 1 Akers, R.F., Lovinger, D.M., Colley, P.A., Linden, D.J. and Routtenberg, A., Translation of protein kinase C activity may mediate hippocampal long-term potentiation, *Science*, 231 (1986) 587-589.
- 2 Albus, J.S., A theory of cerebellar function, *Math. Biosci.*, 10 (1971) 25-61.
- 3 Armstrong, J.N., McIntyre, D.C., Neubort, S. and Soviter, R.S., Granule cell loss following adrenalectomy has no effect on the acquisition, retention and reversal of place learning in the Morris water maze, *Soc. Neurosci. Abstr.*, 17 (1991) 54.5, p. 485.
- 4 Ashton, A.B., Bitgood, S.C. and Moore, J.W., Auditory differential conditioning of the rabbit nictitating membrane response: III. Effects of US shock intensity and duration, *Psychon. Sci.*, 15 (1969) 127-128.
- 5 Baloch, A.A. and Waxman, A.M., Visual learning, adaptive expectations and behavioral conditioning of the mobile robot MAVIN, *Neural Netw.*, 4 (1991) 271-302.
- 6 Baloch, A.A. and Waxman, A.M., Behavioral conditioning of the mobile robot MAVIN. In P. Antognetti and V. Milutinovic (Eds.), *Neural Networks: Concepts, Applications and Implementations, Vol. IV*, Prentice-Hall, New York, 1991.
- 7 Banquet, J.P. and Grossberg, S., Probing cognitive processes through the structure of event-related potentials during learning: an experimental and theoretical analysis, *Appl. Opt.*, 26 (1987) 4931-4946.
- 8 Barto, A.G., Sutton, R.S. and Anderson, C.W., Neuron-like adaptive elements that can solve difficult learning problems, *IEEE Trans.*, SMC-13 (1983) 834-846.
- 9 Berger, T.W., Berry, S.D. and Thompson, R.F., Role of the hippocampus in classical conditioning of aversive and appetitive behaviors. In R.L. Isaacson and K.H. Pribram (Eds.), *The Hippocampus, Vol. 4*, Plenum, New York, 1986, pp. 203-239.
- 10 Berger, T.W., Laham, R.I. and Thompson, R.F., Hippocampal unit-behavior correlations during classical conditioning, *Brain Res.*, 193 (1980) 229-248.
- 11 Berger, T.W. and Orr, W.B., Hippocampectomy selectively disrupts discrimination reversal conditioning of the rabbit nictitating membrane response, *Behav. Brain Res.*, 8 (1983) 49-68.
- 12 Berger, T.W. and Thompson, R.F., Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response. I. The hippocampus, *Brain Res.*, 145 (1978) 323-346.
- 13 Berger, T.W. and Weiss, D.J., Single unit analysis of hippocampal pyramidal cells and their role in classical conditioning of the rabbit nictitating membrane response. In I. Gormezano, W.F. Prokasy and R.F. Thompson (Eds.), *Classical Conditioning*, Lawrence Erlbaum, Hillsdale, NJ, 1987.
- 14 Berry, S.D. and Oliver, C.G., Hippocampal activity during appetitive classical conditioning in rabbits, *Soc. Neurosci. Abstr.*, 8 (1982) 315.
- 15 Bindra, D., Neurophysiological interpretation of the effects of drive and incentive-motivation on general activity and instrumental behavior, *Psychol. Rev.*, 75 (1968) 1-22.
- 16 Boe, E.E., Effect of punishment duration and intensity on the extinction of an instrumental response, *J. Exp. Psychol.*, 72 (1966) 125-131.
- 17 Bok, S.T., *Histology of the Cerebral Cortex*, Elsevier, Amsterdam, 1959, pp. 81 and 391.
- 18 Bolles, R.C. (1967), *Theory of Motivation*, Harper and Row, New York, 1967.
- 19 Borozci, G., Storms, L.H. and Broen, W.E., Response suppression and recovery of responding at different deprivation levels as functions of intensity and duration of punishment, *J. Comp. Physiol. Psychol.*, 58 (1964) 456-459.
- 20 Boss, B.D., Peterson, G.M. and Cowan, W.M., On the numbers of neurons in the dentate gyrus of the rat, *Brain Res.*, 338 (1985) 144-150.
- 21 Boss, B.D., Turlejski, N., Stanfield, B.B. and Cowan, W.M., On the numbers of neurons in fields CA1 and CA3 of the hippocampus of Sprague-Dawley and Wistar rats, *Brain Res.*, 406 (1987) 280-287.
- 22 Bower, G.H., Mood and memory, *Am. Psychol.*, 36 (1981) 129-148.
- 23 Bower, G.H., Gilligan, S.G. and Monteiro, K.P., Selectivity of learning caused by adaptive states, *J. Exp. Psychol. Gen.*, 110 (1981) 451-473.
- 24 Brindley, G.S., The use made by the cerebellum of the information that it receives from sense organs, *Int. Brain Res. Organ. Bull.*, 3 (1964) 80.
- 25 Brown, T.H., Chang, V.C., Ganong, A.H., Keenan, C.L. and Kelso, S.R., Biophysical properties of dendrites and spines that may control the induction and expression of long-term synaptic potentiation. In P.W. Landfield and S.A. Deadwyler (Eds.), *Long-Term Potentiation: From Biophysics to Behavior*, Alan R. Liss, New York, 1988, pp. 201-264.
- 26 Brown, T.H. and Zador, A.M., Hippocampus. In G.M. Shepherd (Ed.), *The Synaptic Organization of the Brain, 3rd edn.*, Oxford University Press, New York, 1990, pp. 346-388.
- 27 Bullock, D. and Grossberg, S., Adaptive neural networks for control of movement trajectories invariant under speed and force rescaling, *Hum. Mov. Sci.*, 10 (1991) 1-51.
- 28 Buonomano, D.V., Baxter, D.A. and Byrne, J.H., Small networks of empirically derived adaptive elements simulate some higher-order features of classical conditioning, *Neural Netw.*, 3 (1990) 507-524.
- 29 Butters, N. and Cermak, L.S., Some analyses of amnesic syndromes in brain-damaged patients. In R.L. Isaacson and K.H. Pribram (Eds.), *The Hippocampus, Vol. 2: Neurophysiology and Behaviour*, Plenum, New York, pp. 377-409.
- 30 Carpenter, G.A. and Grossberg, S., Adaptation and transmitter gating in vertebrate photoreceptors, *J. Theor. Neurobiol.*, 1 (1981) 1-42.
- 31 Carpenter, G.A. and Grossberg, S., A massively parallel architecture for a self-organizing neural pattern recognition machine, *Comput. Vision Graph. Image Process.*, 37 (1987) 54-115.
- 32 Carpenter, G.A. and Grossberg, S., ART 2: self-organization of stable category recognition codes for analog input patterns, *Appl. Opt.*, 26 (1987) 4919-4930.
- 33 Carpenter, G.A. and Grossberg, S., The ART of adaptive pattern recognition by a self-organizing neural network, *Computer*, 21 (1988) 77-88.
- 34 Carpenter, G.A. and Grossberg, S., Neural dynamics of category learning and recognition: attention, memory consolidation and amnesia. In J.L. Davis, R.W. Newburgh, and E.J. Wegman (Eds.), *Brain Structure, Learning, and Memory*. Westview, Boulder, CO, 1988, pp. 233-290. Reprinted in S. Grossberg (Ed.), *The Adaptive Brain, I: Cognition, Learning, Reinforcement, and Rhythm*, Elsevier/North-Holland, Amsterdam, 1987, pp. 238-284.
- 35 Carpenter, G.A. and Grossberg, S., ART 3: hierarchical search

- using chemical transmitters in self-organizing pattern recognition architectures, *Neural Netw.*, 3 (1990) 129–152.
- 36 Carpenter, G.A. and Grossberg, S. (Eds.), *Pattern Recognition by Self-Organizing Neural Networks*, MIT, Cambridge, MA, 1991.
 - 37 Carpenter, G.A. and Grossberg, S., Distributed hypothesis testing, attention shifts, and transmitter dynamics during the self-organization of brain recognition codes. In P. Schuster and W. Singer (Eds.), *Nonlinear Dynamics and Neuronal Networks*, Springer, New York, 1991.
 - 38 Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H. and Rosen, D.B., Fuzzy ARTMAP: a neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Trans. Neural Netw.*, in press. (Technical Report CAS/CNS-TR-91-016, Boston University, Boston, MA.)
 - 39 Carpenter, G.A., Grossberg, S. and Reynolds, J., ARTMAP: supervised real-time learning and classification of nonstationary data by a self-organizing neural network, *Neural Netw.*, 4 (1991) 565–588.
 - 40 Church, R.M., Raymond, G.A. and Beauchamp, R.D., Response suppression as a function of intensity and duration of punishment, *J. Comp. Physiol. Psychol.*, 63 (1967) 39–44.
 - 41 Clark, G.A., McCormick, D.A., Lavond, D.G. and Thompson, R.F., Effects of lesions of cerebellar nuclei on conditioned behavioral and hippocampal neuronal responses, *Brain Res.*, 291 (1984) 125–136.
 - 42 Cohen, N.J., Preserved learning capacity in amnesia: Evidence for multiple memory systems. In L. Squire and N. Butters (Eds.), *The Neuropsychology of Memory*, Guilford, New York, 1984, pp. 83–103.
 - 43 Coleman, S.R. and Gormezano, I., Classical conditioning of the rabbit's (*Oryctolagus cuniculus*) nictitating membrane response under symmetrical CS-US interval shifts, *J. Comp. Physiol. Psychol.*, 77 (1971) 447–455.
 - 44 Collingridge, G.L. and Davies, S.N., NMDA receptors and long-term potentiation in the hippocampus. In J.C. Watkins and G.L. Collingridge (Eds.), *The NMDA Receptor*, Oxford University Press, Oxford, 1989, pp. 123–135.
 - 45 Commons, M.L., Grossberg, S. and Staddon, J.E.R. (Eds.), *Neural Network Models of Conditioning and Action*, Erlbaum Associates, Hillsdale, NJ, 1991.
 - 46 Cooper, J.R., Bloom, F.E. and Roth, R.H., *The Biochemical Basis of Neuropharmacology*, 2nd edn., Oxford University Press, New York, 1974.
 - 47 Damasio, A.R., The brain binds entities and events by multi-regional activation from convergence zones, *Neural Comput.*, 1 (1989) 123–132.
 - 48 Davis, R.E., Xu, X. and Klinger, P.D., NMDA receptor antagonist MK-801 inhibits acquisition of classical fear conditioning in goldfish, *Soc. Neurosci. Abstr.*, 17 (1991) 194.1, p. 485.
 - 49 Deadwyler, S.A., West, M.O. and Lynch, G., Activity of dentate granule cells during learning: differentiation of perforant path inputs, *Brain Res.*, 169 (1979) 29–43.
 - 50 Deadwyler, S.A., West, M.O. and Robinson, J.H., Entorhinal and septal inputs differentially control sensory-evoked responses in the rat dentate gyrus, *Science*, 211 (1981) 1181–1183.
 - 51 Diaz-Granados, J.L., Greene, P.L., Lobaugh, N.J. and Amsel, A., Deficit in response suppression in the juvenile rat after neonatal hippocampal X-irradiation, *Soc. Neurosci. Abstr.*, 17 (1991) 54.5, p. 133.
 - 52 Ebner, T.J. and Bloedel, J.R., Correlation between activity of Purkinje cells and its modification by natural peripheral stimuli, *J. Neurophysiol.*, 45 (1981) 948–961.
 - 53 Eccles, J.C., Forward. In J.C. Watkins and G.L. Collingridge (Eds.), *The NMDA Receptor*, Oxford University Press, New York, pp. ix–xi.
 - 54 Eichenbaum, H. and Cohen, N.J., Representation in the hippocampus: what do hippocampal neurons code?, *Trends Neurosci.*, 11 (1988) 244–248.
 - 55 Errington, M.L., Lynch, M.A. and Bliss, T.V.P., Long-term potentiation in the dentate gyrus induction and increased glutamate release are blocked by d(-)-aminophosphonovalerate, *Neuroscience*, 20 (1987) 279–284.
 - 56 Estes, W.K., Outline of a theory of punishment. In B.A. Campbell and R.M. Church (Eds.), *Punishment and Aversive Behavior*, Appleton-Century-Crofts, New York, 1969.
 - 57 Fujita, M., Adaptive filter model of the cerebellum, *Biol. Cybern.*, 45 (1982) 195–206.
 - 58 Fujita, M., Simulation of adaptive modification of the vestibulo-ocular reflex with an adaptive filter model of the cerebellum, *Biol. Cybern.*, 45 (1982) 207–214.
 - 59 Gaffan, D., Hippocampus: memory, habit and voluntary movement, *Philos. Trans. R. Soc. London, Ser. B*, 308 (1985) 87–99.
 - 60 Gibbon, J., Scalar expectancy and Weber's law in animal timing, *Psychol. Rev.*, 84 (1977) 279–325.
 - 61 Gibbon, J., The origins of scalar timing, *Learn. Motiv.*, 22 (1991) 3–38.
 - 62 Gilbert, P.F.C. and Thach, W.T., Purkinje cell activity during motor learning, *Brain Res.*, 128 (1977) 309–328.
 - 63 Gormezano, I. and Kehoe, E.J., Associative transfer in classical conditioning to serial compounds. In M.L. Commons, R.L. Herrnstein and A.R. Wagner (Eds.), *Quantitative Analyses of Behavior, Vol. 3, Acquisition*, Ballinger, Cambridge, MA, 1984, pp. 297–322.
 - 64 Graf, P., Squire, L.R. and Mandler, G., The information that amnesic patients do not forget, *J. Exp. Psychol. Learn. Mem. Cogn.*, 10 (1984) 164–178.
 - 65 Gray, J.A., *The Neuropsychology of Anxiety: An Enquiry into the Functions of the Septo-Hippocampal System*, Oxford University Press, Oxford, 1982.
 - 66 Grossberg, S., *The Theory of Embedding Fields with Applications to Psychology and Neurophysiology*, Rockefeller Institute for Medical Research, New York, 1964.
 - 67 Grossberg, S., Some physiological and biochemical consequences of psychological postulates, *Proc. Nat. Acad. Sci.*, 60 (1968) 758–765.
 - 68 Grossberg, S., On the production and release of chemical transmitters and related topics in cellular control, *J. Theor. Biol.*, 22 (1969) 325–364.
 - 69 Grossberg, S., On learning and energy-entropy dependence in recurrent and nonrecurrent signed networks, *J. Stat. Phys.*, 1 (1969) 319–350.
 - 70 Grossberg, S., On learning of spatiotemporal patterns by networks with ordered sensory and motor components, I: Excitatory components of the cerebellum, *Stud. Appl. Math.*, 48 (1969) 105–132.
 - 71 Grossberg, S., On the dynamics of operant conditioning, *J. Theor. Biol.*, 33 (1971) 225–255.
 - 72 Grossberg, S., A neural theory of punishment and avoidance, I: Qualitative theory, *Math. Biosci.*, 15 (1972) 39–67.⁷
 - 73 Grossberg, S., A neural theory of punishment and avoidance, II: Quantitative theory, *Math. Biosci.*, 15 (1972) 253–285.
 - 74 Grossberg, S., Neural expectation: cerebellar tretinal analogs of cells fired by learnable and unlearned pattern classes, *Kybernetik*, 10 (1972) 49–57.
 - 75 Grossberg, S., Contour enhancement, short-term memory, and constancies in reverberating neural networks, *Stud. Appl. Math.*, 52 (1973) 217–257.
 - 76 Grossberg, S., A neural model of attention, reinforcement, and discrimination learning, *Int. Rev. Neurobiol.*, 18 (1975) 263–327.
 - 77 Grossberg, S., Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors, *Biol. Cybern.*, 23 (1976) 121–134.
 - 78 Grossberg, S., Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions, *Biol. Cybern.*, 23 (1976) 187–202.
 - 79 Grossberg, S., 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, Vol. 5*, Academic Press, New York, 1978, pp. 233–374. Reprinted in S. Grossberg (Ed.), *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control*, Reidel, Boston, MA, 1982.
 - 80 Grossberg, S., How does a brain build a cognitive code?, *Psychol. Rev.*, 87 (1980) 1–51.

- 81 Grossberg, S., *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control*, Reidel, Boston, MA, 1982.
- 82 Grossberg, S., Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory, *Psychol. Rev.*, 89 (1982) 529-572.
- 83 Grossberg, S., Some psychological correlates of a developmental, cognitive and motivational theory, *Ann. N.Y. Acad. Sci.*, 425 (1984) 58-142.
- 84 Grossberg, S., Neuroethology and theoretical neurobiology, *Behav. Brain Sci.*, 7 (1984) 388-390.
- 85 Grossberg, S. (Ed.), *The Adaptive Brain, I: Cognition, Learning, Reinforcement, and Rhythm*, Elsevier/North-Holland, Amsterdam, 1987.
- 86 Grossberg, S. (Ed.), *The Adaptive Brain, II: Vision, Speech, Language, and Motor Control*, North-Holland, Amsterdam, 1987.
- 87 Grossberg, S., Nonlinear neural networks: principles, mechanisms, and architectures, *Neural Netw.*, 1 (1988) 17-61.
- 88 Grossberg, S., A neural network architecture for Pavlovian conditioning: reinforcement, attention, forgetting, timing. In M.L. Commons, S. Grossberg and J.E.R. Staddon (Eds.), *Neural Network Models of Conditioning and Action*, Erlbaum Associates, Hillsdale, NJ, 1991, pp. 69-122.
- 89 Grossberg, S. and Kuperstein, M., *Neural Dynamics of Adaptive Sensory-Motor Control: Ballistic Eye Movements*, North-Holland, Amsterdam, 1986.
- 90 Grossberg, S. and Kuperstein, M., *Neural Dynamics of Sensory-Motor Control: Expanded Edition*, Pergamon, Elmsford, NY, 1989.
- 91 Grossberg, S. and Levine, D.S., Neural dynamics of attentionally modulated Pavlovian conditioning: blocking, inter-stimulus interval, and secondary reinforcement, *Appl. Opt.*, 26 (1987) 5015-5030.
- 92 Grossberg, S. and Merrill, J.W.L., A neural network model of adaptively timed reinforcement learning and hippocampal dynamics, *Soc. Neurosci. Abstr.*, 17 (1991) 485.
- 93 Grossberg, S. and Schmajuk, N.A., Neural dynamics of Pavlovian conditioning: conditioned reinforcement, inhibition, and opponent processing, *Psychobiology*, 15 (1987) 195-240.
- 94 Grossberg, S. and Schmajuk, N.A., Neural dynamics of adaptive timing and temporal discrimination during associative learning, *Neural Netw.*, 2 (1989) 79-102.
- 95 von Helmholtz, H.L.F., *Handbuch der Physiologischen Optik*, Voss, Leipzig, 1866.
- 96 von Helmholtz, H.L.F., *Treatise on Physiological Optics*, Dover, New York, 1962. (J.P.C. Southall, Translator.)
- 97 Henneman, E., Relation between size of neurons and their susceptibility to discharge, *Science*, 26 (1957) 1345-1347.
- 98 Henneman, E., The size-principle: a deterministic output emerges from a set of probabilistic connections, *J. Exp. Biol.*, 115 (1985) 105-112.
- 99 Hindmarsh, A.C., ODEPACK: systematized collection of ODE solvers. In R.S. Stepleman et al. (Eds.), *Scientific Computing*, North-Holland, Amsterdam, 1983, pp. 55-64.
- 100 Hoehler, F.K. and Thompson, R.F., Effects of the interstimulus (CS-US) interval on hippocampal unit activity during classical conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*), *J. Comp. Physiol. Psychol.*, 94 (1980) 201-215.
- 101 Holder, M.D. and Roberts, S., Comparison of timing and classical conditioning, *J. Exp. Psychol. Anim. Behav. Process.*, 11 (1985) 172-193.
- 102 Houk, J.C., Singh, S.P., Fisher, C. and Barto, A.G., An adaptive sensorimotor network inspired by the anatomy and physiology of the cerebellum. In W.T. Miller, R.S. Sutton, and P.J. Werbos (Eds.), *Neural Networks for Control*, MIT Press, Cambridge, MA, 1989, pp. 301-348.
- 103 Isaacson, R.L. and Pribram, K.H. (Eds.), *The Hippocampus, Vol. 4*, Plenum, New York, 1986.
- 104 Ito, M., The control mechanism of cerebellar motor systems. In F.O. Schmidt and F.G. Worden (Eds.), *The Neurosciences Third Study Program*, MIT Press, Cambridge, MA, 1974.
- 105 Ito, M., *The Cerebellum and Neural Control*. Raven, New York, 1984.
- 106 James, W., *Principles of Psychology*, Henry Holt, New York, 1890.
- 107 James, W., *Pragmatism*, Longmans and Green, New York, 1914.
- 108 John, E.R., *Mechanisms of Memory*, Academic Press, New York, 1967.
- 109 Jones, R. and Keck, M.J., Visual evoked response as a function of grating spatial frequency, *Invest. Ophthalmol. Vis. Sci.*, 17 (1978) 652-659.
- 110 Kamin, L.J., Predictability, surprise, attention, and conditioning. In B.A. Campbell and R.M. Church (Eds.), *Punishment and Aversive Behavior*, Appleton-Century-Crofts, New York, 1969.
- 111 Keehn, J.D., Effect of shock duration on Sidman avoidance response rates, *Psychol. Rep.*, 13 (1963) 852.
- 112 Kehoe, E.J., Marshall-Goodell, B. and Gormezano, I., Differential conditioning of the rabbit's nictitating membrane response to serial compound stimuli, *J. Exp. Psychol. Anim. Behav. Process.*, 13 (1987) 17-30.
- 113 Kettner, R.E. and Thompson, R.F., Auditory signal detection and decision processes in the nervous system, *J. Comp. Physiol. Psychol.*, 96 (1982) 328-331.
- 114 Konorski, J., *Integrative Activity of the Brain*, University of Chicago, 1967.
- 115 LeDoux, J.E., Cognitive-emotional interactions in the brain, *Cogn. Emot.*, 3 (1989) 267-289.
- 116 LeDoux, J.E., Iwata, J., Cichetti, P. and Reis, D.J., Differential projections of the central amygdaloid nucleus mediate autonomic and behavioral correlates of conditioned fear, *J. Neurosci.*, 8 (1988) 2517-2529.
- 117 LeDoux, J.E., Romanski, L. and Xagoraris, A., Indelibility of subcortical emotional memories, *J. Cogn. Neurosci.*, 1 (1989) 238-243.
- 118 Levine, D.S. and Prueitt, P.S., Modeling some effects of frontal lobe damage—novelty and perseveration, *Neural Netw.*, 2 (1989) 103-116.
- 119 Levine, D.S. and Prueitt, P.S., Simulations of conditioned perseveration and novelty preference from frontal lobe damage. In M.L. Commons, S. Grossberg and J.E.R. Staddon (Eds.), *Neural Network Models of Conditioning and Action*, Erlbaum Associates, Hillsdale, NJ, 1991, pp. 123-148.
- 120 Levy, W.B., Brassel, S.E. and Moore, S.D., Partial quantification of the associative synaptic learning rule of the dentate gyrus, *Neuroscience*, 81 (1983) 799-808.
- 121 Levy, W.B. and Desmond, N.L., The rules of elemental synaptic plasticity. In W.B. Levy, J. Anderson and S. Lehmkuhle (Eds.), *Synaptic Modification, Neuron Selectivity and Nervous System Organization*, Erlbaum Associates, Hillsdale, NJ, 1985, pp. 105-121.
- 122 Lynch, G., McGaugh, J.L. and Weinberger, N.M. (Eds.), *Neurobiology of Learning and Memory*, Guilford, New York, 1984.
- 123 Maricq, A.V., Roberts, S. and Church, R.M., Methamphetamine and time estimation, *J. Exp. Psychol. Anim. Behav. Process.*, 7 (1981) 18-30.
- 124 Marr, D., A theory of cerebellar cortex, *J. Physiol.*, 202 (1969) 437-470.
- 125 Mattis, S. and Kovner, R., Amnesia is as amnesia does: toward another definition of the anterograde amnesias. In L. Squire and N. Butters (Eds.), *Neuropsychology of Memory*, Guilford, New York, 1984, pp. 115-121.
- 126 McCormick, D.A., Clark, G.A., Lavond, D.G. and Thompson, R.F., Initial localization of the memory trace for a basic form of learning, *Proc. Nat. Acad. Sci. U.S.A.*, 79 (1982) 2731-2735.
- 127 Meck, W.H. and Church, R.M., Cholinergic modulation of the content of temporal memory, *Behav. Neurosci.*, 101 (1987) 457-464.
- 128 Millenson, J.R., Kehoe, E.J. and Gormenzano, I., Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS-US intervals, *Learn. Motiv.*, 8 (1977) 351-366.
- 129 Miller, N.E., Some reflections on the law of effect produce a new alternative to drive reduction. In M.R. Jones (Ed.), *Ne-*

- braska Symposium on Motivation*, University of Nebraska Press, Lincoln, 1963.
- 130 Musselwhite, M.J. and Jeffreys, D.A., The influence of spatial frequency on the reaction times and evoked potentials recorded to grating pattern stimuli, *Vis. Res.*, 25 (1985) 1545-1555.
 - 131 O'Keefe, J. and Nadel, L., *The Hippocampus as a Cognitive Map*, Oxford University Press, Oxford, 1978.
 - 132 Olds, J., *Drives and Reinforcements: Behavioral Studies of Hypothalamic Functions*, Raven, New York, 1977.
 - 133 Optican, L.M. and Robinson, D.A., Cerebellar-dependent adaptive control of primate saccadic system, *J. Neurophysiol.*, 44 (1980) 108-1076.
 - 134 Parker, D.M. and Salzen, E.A., Latency changes in the human visual evoked response to sinusoidal gratings, *Vis. Res.*, 17 (1977) 1201-1204.
 - 135 Parker, D.M. and Salzen, E.A., The spatial selectivity of early and late waves within the human visual evoked response, *Perception*, 6 (1977) 85-95.
 - 136 Parker, D.M., Salzen, E.A. and Lishman, J.R., Visual-evoked responses elicited by the onset and offset of sinusoidal gratings: Latency, waveform, and topographic characteristics, *Invest. Ophthalmol. Vis. Sci.*, 22 (1982) 657-680.
 - 137 Parker, D.M., Salzen, E.A. and Lishman, J.R., The early waves of the visual evoked potential to sinusoidal gratings: Responses to quadrant stimulation as a function of spatial frequency, *Electroencephalogr. Clin. Neurophysiol.*, 53 (1982) 427-435.
 - 138 Pavlov, I.P., *Conditioned Reflexes*, Constable and Company, London, 1927. (Reprinted by Dover Publications, 1960.)
 - 139 Petzold, L.R., Automatic selection of methods for solving stiff and nonstiff systems of ordinary differential equations, *SIAM J. Sci. Stat. Comput.*, 4 (1983) 136-148.
 - 140 Plant, G.T., Zimmern, R.L. and Durden, K., Transient visually evoked potentials to the pattern reversal and onset of sinusoidal gratings, *Electroencephalogr. Clin. Neurophysiol.*, 56 (1983) 147-158.
 - 141 Port, R.L., Mikhail, A.A. and Patterson, M.M., Differential effects of hippocampectomy on classically conditioned rabbit nictitating membrane response related to interstimulus interval, *Behav. Neurosci.*, 99 (1985) 200-208.
 - 142 Pribram, K.H., The hippocampal system and recombinant processing. In R.L. Isaacson and K.H. Pribram (Eds.), *The Hippocampus, Vol. 4*, Plenum, New York, 1986, pp. 329-370.
 - 143 Rickert, E.J., Bennett, T.L., Lane, P.L. and French, J., Hippocampectomy and the attenuation of blocking, *Behav. Biol.*, 22 (1978) 147-160.
 - 144 Roberts, S., Isolation of an internal clock, *J. Exp. Psychol. Anim. Behav. Process.*, 7 (1981) 242-268.
 - 145 Roberts, W.A., Cheng, K. and Cohen, J.S., Timing light and tone signals in pigeons, *J. Exp. Psychol. Anim. Behav. Process.*, 15 (1989) 23-25.
 - 146 Robinson, G.B., Contribution of NMDA-mediated neuronal activity to classical conditioning of the rabbit nictitating membrane and conditioning-related potentiation of perforant path-granule cell responses, *Soc. Neurosci. Abstr.*, 17 (1991) 130.12, p. 485.
 - 147 Ron, S. and Robinson, D.A., Eye movements evoked by cerebellar stimulation in the alert monkey, *J. Neurophysiol.*, 36 (1973) 1004-1021.
 - 148 Rosene, D.L. and van Hoesen, G.W., Hippocampal efferents reach widespread areas of cerebral cortex and amygdala in the rhesus monkey, *Science*, 198 (1977) 315-317.
 - 149 Schmajuk, N.A. and DiCarlo, J.J., Neural dynamics of hippocampal modulation of classical conditioning. In M.L. Commons, S. Grossberg and J.E.R. Staddon (Eds.), *Neural Network Models of Conditioning and Action*, Erlbaum Associates, Hillsdale, NJ, 1991, pp. 149-180.
 - 150 Schmajuk, N.A., Spear, N.E. and Isaacson, R.L., Absence of overshadowing in rats with hippocampal lesions, *Physiol. Psychol.*, 11 (1983) 59-62.
 - 151 Singer, W., Neuronal activity as a shaping factor in the self-organization of neuron assemblies. In: E. Basar, H. Flohr, H. Haken and A.J. Mandell (Eds.), *Synergetics of the Brain*, Springer, New York, 1983.
 - 152 Skrandies, W., Scalp potential fields evoked by grating stimuli: Effects of spatial frequency and orientation, *Electroencephalogr. Clin. Neurophysiol.*, 58 (1984) 325-332.
 - 153 Smith, M.C., CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response, *J. Comp. Physiol. Psychol.*, 3 (1968) 679-687.
 - 154 Sokolov, E.N., *Perception and the Conditioned Reflex*, Moscow University Press, Moscow, 1958.
 - 155 Sokolov, E.N., *Mechanisms of Memory*, Moscow University Press, Moscow, 1968.
 - 156 Solomon, P.R., The role of hippocampus in blocking and conditioned inhibition of the rabbit's nictitating membrane response, *J. Comp. Physiol. Psychol.*, 91 (1977) 407-417.
 - 157 Solomon, P.R., Temporal versus spatial information processing views of hippocampal functions, *Psychol. Bull.*, 86 (1979) 1272-1279.
 - 158 Solomon, P.R., A time and place for everything? Temporal processing views of hippocampal function with special reference to attention, *Physiol. Psychol.*, 8 (1980) 254-261.
 - 159 Solomon, P.R., Vander Schaaf, E., Thompson, R.F. and Weisz, D.J., Hippocampus and trace conditioning of the rabbit's nictitating membrane response, *Soc. Neurosci. Abstr.*, 9 (1983) 645.
 - 160 Squire, L.R. and Butters, N. (Eds.), *Neuropsychology of Memory*, Guilford, New York, 1984.
 - 161 Squire, L.R. and Cohen, N.J., Human memory and amnesia. In G. Lynch, J. McGaugh, and N.M. Weinberger (Eds.), *Neurobiology of Learning and Memory*, Guilford Press, New York, 1984, pp. 3-64.
 - 162 Squire, L.R., Shimamura, A.P. and Amaral, D.G., Memory and the hippocampus. In J.H. Byrne and W.O. Berry (Eds.), *Neural Models of Plasticity: Experimental and Theoretical Approaches*, Academic Press, San Diego, 1989, pp. 208-239.
 - 163 Staddon, J.E.R., *Static and Dynamic Competition*, Handout at Tutorial Conference on Neural Modelling, Scottsdale, AZ, 1983.
 - 164 Strouthes, A., Effect of CS-onset, UCS-termination delay, UCS duration, CS-onset interval and number of CS-UCS pairings on conditioned fear response, *J. Exp. Psychol.*, 69 (1965) 287-291.
 - 165 Sutherland, R.J., Dringenberg, H.C., Hoising, J.M. and Skelton, R.W., Is LTE in the hippocampus necessary for place learning?, *Soc. Neurosci. Abstr.*, 17 (1991) 193.11, p. 483.
 - 166 Thompson, L.T. and Disterhoft, J.F., The NMDA receptor/channel complex and associative learning: opposite effects of PCP and O-cycloserine, *Soc. Neurosci. Abstr.*, 17 (1991) 194.8, p. 486.
 - 167 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., Neuronal substrates of associative learning in the mammalian brain. In D.L. Aldon and J. Farley (Eds.), *Primary Neural Substrates of Learning and Behavioral Change*, Cambridge University Press, New York, 1984, pp. 71-99.
 - 168 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., Neuronal substrates of discrete, defensive conditioned reflexes, conditioned fear states, and their interactions in the rabbit. In I. Gormezano, W.F. Prokasy and R.F. Thompson (Eds.), *Classical Conditioning, 3rd edn.*, Erlbaum Associates, Hillsdale, NJ, 1987, pp. 371-399.
 - 169 Tolman, E.C., *Purposive Behavior in Animals and Men*, Appleton-Century-Crofts, New York, 1932.
 - 170 Valenstein, E.S., Biology of drives. In F.O. Schmitt, T. Melnechuk, G.C. Quarten, and G. Adleman (Eds.), *Neurosciences Research Symposium Summaries, Vol. 3*, MIT Press, Cambridge, MA, 1969.
 - 171 Vassilev, A., Manahilov, V. and Mitov, D., Spatial frequency and pattern onset-offset response, *Vis. Res.*, 23 (1983) 1417-1422.
 - 172 Vassilev, A. and Strashimirov, D., On the latency of human visually evoked response to sinusoidal gratings, *Vis. Res.*, 19 (1979) 843-846.
 - 173 Warrington, E.K. and Eiskrantz, L., The amneic syndrome: consolidation or retrieval?, *Nature*, 228 (1970) 628-630.
 - 174 Warrington, E.K. and Weiskrantz, L., The effect of prior learn-

- ing on subsequent retention in amnesic patients, *Neuropsychologia*, 12 (1974) 419–428.
- 175 Weikard, C. and Berger, T.W., (1986) Hippocampectomy disrupts classical conditioning of cross-modality discrimination reversal learning of the rabbit nictitating membrane response, *Behav. Brain Res.*, 22 (1986) 85–89.
- 176 Wilkie, D.M., Stimulus intensity affects pigeons' timing behavior: implications for an internal clock model, *Anim. Learn. Behav.*, 15 (1987) 35–39.
- 177 Williamson, S.J., Kaufman, I. and Brenner, D., Latency of the neuromagnetic response of the human visual cortex, *Vis. Res.*, 18 (1978) 107–110.
- 178 Zola-Morgan, S.M. and Squire, L.R., The primate hippocampal formation evidence for a time-limited role in memory storage, *Science*, 250 (1990) 288–290.

APPENDIX

Simulation methods. All simulations were performed on an Iris-4D/240 superwork-station using double pre-

cision representations of all values. The Iris-4D series is based upon a microcomputer that conforms to the IEEE floating point standard for accuracy. Simulations were performed by integrating the dynamical system that defines the model. Integration of this system was performed using LSODA (The Livermore Solver for Ordinary Differential equations with Automatic method-switching for stiff and non-stiff systems)^{99,139}.

The time scale was chosen to be consistent with that used in Grossberg and Schmajuk⁹⁴. Trials were set to be 2 'units' long. In ref. 94, trials were set to be 2,000 'milliseconds' long. All parameters from their model were preserved as they appeared in the original, except that they were multiplied by the scale factor 1,000.