

ORIGINAL CONTRIBUTION

Neural Dynamics of Adaptive Timing and Temporal Discrimination During Associative Learning

STEPHEN GROSSBERG† AND NESTOR A. SCHMAJUK‡

Center for Adaptive Systems, Boston University

(Received April 1988; revised and accepted 19 October 1988)

Abstract—A neural network model that controls behavioral timing is described and simulated. This model, called the Spectral Timing Model, controls a type of timing whereby an animal or robot can learn to wait for an expected goal by discounting expected nonoccurrences of a goal object until the expected time of arrival of the goal. If the goal object does not then materialize, the animal can respond to unexpected nonoccurrences of the goal with appropriate changes in information processing and exploratory behavior. The model is a variant of the gated dipole model of opponent processing. When the gated dipole model is generalized to include a spectrum of cellular response rates within a large population of cells, the model's total output signal generates accurate learned timing properties that collectively provide a good quantitative fit to animal learning data. In particular, the Spectral Timing Model utilizes the habituating transmitter gates and adaptive long-term memory traces that are characteristic of gated dipole models. The Spectral Timing Model is embedded into an Adaptive Resonance Theory (ART) neural architecture for the learning of correlations between internal representations of recognition codes and reinforcement codes. This type of learning is called conditioned reinforcer learning. The two types of internal representations are called sensory representations (*S*) and drive representations (*D*). Activation of a drive representation *D* by the Spectral Timing Model inhibits output signals from the orienting subsystem (*A*) of the ART architecture and activates a motor response. The inhibitory pathway helps to prevent spurious resets of short-term memory, forgetting, and orienting responses from being caused by events other than the goal object prior to the expected arrival time of the goal. Simulated data properties include the inverted *U* in learning as a function of the interstimulus interval (*ISI*) that occurs between onset of the conditioned stimulus (*CS*) and the unconditioned stimulus (*US*); correlations of peak time, standard deviation, Weber fraction, and peak amplitude of the conditioned response as a function of the *ISI*; increase of conditioned response amplitude, but not its timing, with *US* intensity; speed-up of the timing circuit by an increase in *CS* intensity or by drugs that increase concentrations of brain dopamine or acetylcholine; multiple timing peaks in response to learning conditions using multiple *ISIs*; and conditioned timing of cell activation within the hippocampus and of the contingent negative variation (*CNV*) event-related potential. The results on speed-up by drugs that increase brain concentrations of dopamine and acetylcholine support a 1972 prediction that the gated dipole habituating transmitter is a catecholamine and its long-term memory trace transmitter is acetylcholine. It is noted that the timing circuit described herein is only one of several functionally distinct neural circuits for governing different types of timed behavior competence.

Keywords—Neural networks, Associative learning, Robotics, Classical conditioning, Timing, Reinforcement, Adaptive resonance theory, Cognition, Hippocampus, Nictitating membrane, Orienting response.

†Supported in part by the Air Force Office of Scientific Research (AFOSR F49620-86-C-0037 and AFOSR F49620-87-C-0018) and the National Science Foundation (NSF IRI-84-17756).

‡Supported in part by the National Science Foundation (NSF IRI-84-17756). The present address of Professor Schmajuk is: Northwestern University, Psychology Department, Evanston, IL 60201.

Acknowledgements: The authors thank Carol Yanakakis and Cynthia Suchta for their valuable assistance in the preparation of the manuscript and illustrations.

Requests for reprints should be sent to Stephen Grossberg, Center for Adaptive Systems, Boston University, Mathematics Department, 111 Cummington Street, Boston, MA 02215.

1. INTRODUCTION: TIMING THE EXPECTED DELAY OF A GOAL OBJECT IN A SPATIALLY DISTRIBUTED AND NONSTATIONARY WORLD

This article presents a model of a neural circuit that controls behavioral timing. There are several different types of brain processes that organize the temporal unfolding of serial order in behavior. The present article describes in detail a model of one type of timing circuit, and outlines how this circuit may be embedded in larger neural systems that regulate

several different types of temporal organization. It seems to us that such timing circuits are just as important for the survival of animals as they are for the design of robots that are capable of freely moving in a spatially distributed world that is characterized by unexpected events and nonstationary statistics.

Many goal objects in such a world may be delayed subsequent to the actions that elicit them, or the environmental events that signal their subsequent arrival. Were all causes followed immediately by their consequences, the world would be a much simpler place to live. In the world as it is, humans and many animal species can learn to wait for the anticipated arrival of a delayed goal object. In part, this capability enhances the efficiency of the consummatory behavior that is triggered by the arrival of the goal object, such as eating when the goal object is food, because the animal can time the preparations to eat so that they are synchronized with the arrival of the food.

The need for behavioral timing becomes even more important in the lives of animals that are capable of exploring their environments for novel sources of gratification. Although the evolution of efficient locomotion greatly enhanced the range of alternative goals that an animal could sample, it also created the danger that the animal may never be able to consummate at all. For example, if an animal could not inhibit its exploratory behavior, then it could easily starve to death by restlessly moving from place to place, unable to remain in one place long enough to carry out the consummatory behaviors needed to acquire food there. On the other hand, if an animal inhibited its exploratory behavior for too long, and remained in one place waiting for an expected source of food to materialize, then it could starve to death if food was not, after all, forthcoming.

2. TIMING THE BALANCE BETWEEN EXPLORATION FOR NOVEL REWARDS AND CONSUMMATION OF EXPECTED REWARDS

Thus the animal's task is to accurately time the *expected* delay of a goal object based upon its previous experiences in a given situation. It needs to regulate the balance between its exploratory behavior aimed at searching for novel sources of reward, and its consummatory behavior aimed at acquiring 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.

3. DISTINGUISHING EXPECTED NONOCCURRENCES FROM UNEXPECTED NONOCCURRENCES: INHIBITING THE NEGATIVE CONSEQUENCES OF EXPECTED NONOCCURRENCES

The type of timed behavior described above is restricted to calibrating the delay of a single behavioral act, rather than organizing a correctly timed and speed-controlled sequence of acts. The key problem that needs to be mechanistically understood is illustrated by the following example. Suppose that, after pushing a lever, an animal typically receives a food pellet from a food magazine two seconds later. Suppose that the animal orients to the food magazine right after pushing the lever. When the animal inspects the food magazine, it perceives the nonoccurrence of food during the subsequent two seconds. These nonoccurrences disconfirm the sensory expectation that food will appear in the magazine. Moreover, the perceptual processing cycle that processes this sensory information occurs at a much faster rate than two seconds, so that it can compute this sensory disconfirmation many times before the two second delay has elapsed.

The key issue is: What spares the animal from erroneously reacting to these *expected nonoccurrences* of food during the first two seconds as predictive failures? Why does not the animal immediately become frustrated by the nonoccurrence of food and release exploratory behavior aimed at searching for food in another place? On the other hand, if food does not appear after two seconds have elapsed, why does the animal then react to the *unexpected nonoccurrence* of food by becoming frustrated and releasing exploratory behavior?

We assert that a primary role of the timing mechanism is to inhibit, or *gate*, the process whereby sensory mismatches trigger the orienting and reinforcing mechanisms that would otherwise reset the animal's attentional focus, negatively reinforce its previous consummatory behavior, and release its exploratory behavior. The process of *registering* these sensory mismatches or matches, as the case might be, is not inhibited. Indeed, if the food happened to appear earlier than expected, the animal could certainly perceive its occurrence and begin to respond accordingly. The sensory matching process, as such, is thus not inhibited by the timing mechanism. Rather, the effects of sensory mismatches upon processes of sensory reset and reinforcement are inhibited.

This inhibitory action is assumed to be part of a more general competition that occurs between the motivational, or arousal, sources that energize different types of behavior. Exploratory behaviors enable the animal to come into contact with novel goal objects. Such behaviors are assumed to be energized by endogenously active motivational sources. Hence,

unless they are actively inhibited, these endogenously active arousal sources could remove the animal from all sources of delayed reward. Consummatory behaviors, such as eating, enable the animal to complete behavioral cycles involving familiar and accessible goal objects. The inhibitory action posited above is from the motivational sources of consummatory behaviors to the motivational sources of orienting and exploratory behaviors.

It is also assumed that the consummatory arousal sources are in mutual competition, enabling only the strongest combinations of sensory, reinforcing, and homeostatic signals to control observable behaviors (Grossberg, 1982, Chap. 6; Staddon, 1983). Thus the posited competition is a special case of the general hypothesis that the output signals from all motivational sources compete for the control of observable behaviors.

To explain how this process works, the present article is organized into two parts:

Part I describes a model of the timing circuit and shows that it can be used to quantitatively explain data from a number of classical and instrumental conditioning experiments about how timed behavior is learned.

Part II shows how this timing circuit can be embedded in a larger neural system to carry out the gating function described above. This larger system is a specialized Adaptive Resonance Theory, or ART, circuit that has been progressively developed in a number of articles since its first appearances in Grossberg (1975, 1978). These and relevant subsequent articles are brought together in several books (Grossberg, 1982, 1987, 1988). The present article provides a summary of the major circuit concepts.

PART I SPECTRAL TIMING MODEL

4. SPECTRAL TIMING MODEL: AN APPLICATION OF GATED DIPOLE THEORY

The timing model presented herein grew out of, and forms part of, a larger theory of cognitive-emotional interactions (Grossberg, 1982, 1987, 1988; Grossberg & Levine, 1987; Grossberg & Schmajuk, 1987). These are the interactions whereby reinforcing events influence the course of conditioning or associative learning through time and thereby regulate the salience of the events to which an animal will subsequently attend. The model is evaluated by demonstrating its competence in explaining data about how animals time their responses during conditioning experiments.

The two major experimentally controlled events during a conditioning experiment are the conditioned

stimulus (CS) and the unconditioned stimulus (US). The CS is a sensory stimulus which does not initially possess the reinforcing properties of the US, but gains (some of) these properties by being paired with the US during learning trials. We denote by $I_{CS}(t)$ the internal input generated by the CS to the timing circuit, and by $I_{US}(t)$ the internal input generated by the US to the timing circuit.

The timing model specializes a design for an opponent processing network, called a *gated dipole*, that was introduced in Grossberg (1972a, 1972b). One version of the model is described. It is called the Spectral Timing Model for reasons described below. The model developed herein uses only feedforward anatomical pathways. On the other hand, as is often the case, the learning is controlled by feedback signals within these pathways.

The circuit diagram of the Spectral Timing model is schematized in Figure 1. A key property of the model is that the CS activates a population of cells whose members react at different rates, according to a spectrum of rates α_i . Neural populations whose elements are distributed along a temporal or spatial parameter are familiar throughout the nervous system. Two examples are the *size principle*, which governs variable rates of responding in spinal motor centers (Henneman, 1957, 1985), and the spatial frequency-tuned cells of the visual cortex, which also react at different rates (Jones & Keck, 1978; Musselwhite & Jeffreys, 1985; Parker & Salzen, 1977a, 1977b; Parker, Salzen, & Lishman 1982a, 1982b; Plant, Zimmern, & Durden, 1983; Skrandies, 1984; Vassilev & Strashimirov, 1979; Vassilev, Manahilov, & Mitov, 1983; Williamson, Kaufman, & Brenner, 1978).

SPECTRAL TIMING EQUATIONS

Spectral Activation

$$\frac{d}{dt} x_i = \alpha_i [-Ax_i + (1 - Bx_i)I_{CS}(t)]; \quad (1)$$

Transmitter Gate

$$\frac{d}{dt} y_i = C(1 - y_i) - Df(x_i)y_i, \quad (2)$$

where $f(x_i)$ is a sigmoid signal function of the form

$$f(x_i) = \frac{x_i^n}{\beta^n + x_i^n}; \quad (3)$$

Associative Learning (LTM Trace)

$$\frac{d}{dt} z_i = Ef(x_i)y_i[-z_i + I_{US}(t)]; \quad (4)$$




FIGURE 1. Circuit diagram of the Spectral Timing Model. The function $I_{CS}(t)$ denotes a step function input that is proportional to the CS intensity and stays on after the CS offset; x_i denote cell activities with different growth rates α_i ; z_i denote adaptive long-term memory traces; and $R(t)$ denotes the total circuit output. In version (a) of the model, the z_i are computed in terminals of the presynaptic pathways converging upon the output neuron, and the I_{US} activates them presynaptically. In (b), the z_i are computed as part of the postsynaptic membrane of the output neuron, and the I_{US} activates them via a postsynaptic route.

Output Signal

$$R = \left[\sum_i f(x_i)y_i z_i - F \right] \quad (5)$$

where

$$[w]^+ = \begin{cases} w & \text{if } w > 0 \\ 0 & \text{if } w \leq 0. \end{cases} \quad (6)$$

A. The Activation Spectrum

The function $I_{CS}(t)$ is assumed to be a step function whose amplitude is proportional to the CS intensity, and which stays on for a fixed time after CS offset because it is internally stored in short-term memory

(STM). Figure 2 depicts a typical relationship between CS, $I_{CS}(t)$, and the US input $I_{US}(t)$. Input $I_{CS}(t)$ activates all potentials x_i in (1) of the cells in its target population. The potentials x_i respond at rates proportional to α_i , $i = 1, 2, \dots, n$.

Each potential x_i generates the output signal $f(x_i)$. Figure 3a depicts the results of a computer simulation in which $f(x_i(t))$ is plotted as a function of time t for values of α_i ranging from .2 ("fast cells") to .0025 ("slow cells").

B. The Habituation Spectrum

Each output signal $f(x_i)$ activates a neurotransmitter y_i . According to equation (2), process y_i accumulates




FIGURE 2. Temporal arrangement of a 50-ms CS and a 50-ms US separated by a 100-ms ISI. I_{CS} is the step function activated by the CS that inputs to the Spectral Timing Model.

to a constant target level 1, via term $C(1 - y_i)$, and is inactivated, or *habituates*, due to a mass action interaction with signal $f(x_i)$, via term $-Df(x_i)y_i$. Although the rate parameters C and D that govern each process y_i are independent of i , the different rates α_i at which each x_i is activated causes the corresponding y_i to become habituated at a different rate. A habituation spectrum is thereby generated at which the y_i processes are successively inactivated. The signal functions $f(x_i(t))$ in Figure 3a generate the habituation spectrum of $y_i(t)$ curves shown in Figure 3b.

C. The Gated Signal Spectrum

Each signal $f(x_i)$ interacts with y_i via mass action. This process is also called the *gating* of $f(x_i)$ by y_i

to yield a net signal g_i proportional to $f(x_i)y_i$. Each of these gated signals, as a function of time $g_i(t) \equiv f(x_i(t))y_i(t)$, has a different rate of growth and decay. The set of all these curves thereby generates a gated signal spectrum, which is shown in Figure 3c. The curves in Figure 3c exhibit the following properties:

1. Each function $g_i(t)$ is a unimodal function of time, where function $g_i(t)$ achieves its maximum value M_i at time T_i .
2. T_i is an increasing function of i ; and
3. M_i is a decreasing function of i .

D. Temporally Selective Associative Learning

Each *long-term memory (LTM) trace* z_i in (4) is activated by its own temporally selective sampling signal g_i . The sampling signal g_i turns on the learning process, and causes z_i to approach I_{US} during the sampling interval at a rate proportional to g_i . Each z_i thus grows by an amount that reflects the degree to which the curves $g_i(t)$ and $I_{US}(t)$ have simultaneously large values through time.

The time interval between CS onset and US onset is called the *interstimulus interval*, or ISI. The individual LTM traces differ in their ability to learn at different values of the ISI. This is the basis of the network's timing properties.

Figure 4 illustrates how six different LTM traces $z_i, i = 1, \dots, 6$, learn during a simulated learning




FIGURE 3. The spectrum of reactions to a step input I_{CS} : (a) Eighty signal functions $f(x_i(t)), i = 1, 2, \dots, 80$, are plotted as a function of t . (b) The corresponding eighty habituating transmitter gates $y_i(t)$ are plotted as a function of t . (c) The corresponding gated signals $g_i(t) = f(x_i(t))y_i(t)$ are plotted as a function of t . Parameters are: $\alpha_i = .2i^{-1}$ for $i = 1, 2, \dots, 80$; $A = 1$; $B = 1$; $C = .0001$; $D = .125$; $\beta = .8$; $n = 8$; $I_{CS}(t) = 1$ for $t > 0$. In all simulations, one time step represents 1 ms and all $f(x_i(0)) = 0$ and $y_i(0) = 1$.




FIGURE 4. Selective learning within different spectral populations at a fixed ISI = 500 ms. Each three-image panel from (a) to (f) represents the gated signal $g_i(t)$ [top], long-term memory trace $z_i(t)$ [middle], and doubly gated signal $h_i(t) = g_i(t)z_i(t)$ [bottom], at a different value of i . In (a), $i = 1$; in (b), $i = 10$; in (c), $i = 20$; in (d), $i = 30$; in (e), $i = 40$; in (f), $i = 50$. The same parameters as in Figure 3 were used. In addition, $E = .01$ and $I_{us}(t) = 10$ for $t \in (500, 550)$ and $= 0$ otherwise.

experiment. The CS and US are paired during 4 learning trials, after which the CS is presented alone on a single performance trial.

E. The Doubly Gated Signal Spectrum

The CS input $I_{CS}(t)$ remains on and constant throughout the duration of each learning trial. The US input $I_{US}(t)$ is presented after an ISI of 500 ms and then remains on for 50 ms. The upper panel in each part of the figure depicts the gated signal function $g_i(t)$ with α_i chosen at progressively slower rates. The middle panel plots the corresponding LTM trace $z_i(t)$, and the lower panel plots the double gated signal $h_i(t) = f(x_i(t))y_i(t)z_i(t)$. Each doubly gated signal function $h_i(t)$ registers how well the timing of CS and US is registered by the i th processing channel. Note that in Figure 4c, whose gated signal $g_i(t)$ peaks at approximately 500 ms the LTM trace $z_i(t)$ exhibits maximum learning. The doubly gated signal $h_i(t)$ also shows a maximal enhancement due to learning, and exhibits peaks of activation at approximately 500 ms after onset of the CS on each trial. This behavior is also generated on the fifth trial, during which only the CS is presented.

F. The Output Signal

The output signal $R(t)$ defined in eqn (5) is the sum of all the doubly gated signal functions $h_i(t)$ minus a threshold F . The output signal computes the cumulative learned reaction of all the cells to the input pattern.

Figure 5a plots the output signal generated in a computer experiment through time across all five trials, using an ISI of 400 ms. In Figure 5b, successive responses in Figure 5a are superimposed to show how they are aligned with respect to the ISI and increase due to learning or successive trials. Figure 5c plots all of the doubly gated signal functions $h_i(t)$ that are summated to form $R(t)$ on the fifth trial. Figure 5d plots all the gated signal functions $g_i(t)$ whose multiplication by $z_i(t)$ generates the $h_i(t)$ curves. Together these Figures illustrate how function $R(t)$ generates an accurately timed response from the cumulative partial learning of all cells in the population spectrum.

5. EFFECT OF INCREASING ISI AND US INTENSITY

Figures 6a–6c plot the curves that are generated by ISIs of 0, 500, and 1000 ms. In every case, the learned cumulative response $R(t)$ is accurately centered at the correct ISI.

Figure 7 plots the functions $R(t)$ that are generated by different ISIs in a series of learning experiments.

There are the $R(t)$ functions generated on the tenth trial of each experiment in response to a CS alone, after four trials of prior learning, with all time axes synchronized with CS onset. In Figure 7a, the $I_{US}(t)$ was chosen twice as large as in Figure 7b. Halving $I_{US}(t)$ amplitude reduces the $R(t)$ amplitudes without changing their timing or overall shape. Note that the envelope of the $R(t)$ functions increases and then decreases through time, and that the individual $R(t)$ functions corresponding to larger ISIs are broader.

6. COMPARISON WITH NICTITATING MEMBRANE CONDITIONING DATA

The computer simulations summarized in Figure 7 are strikingly similar to the data of Smith (1968) summarized in Figure 8. Smith (1968) studied the effect of manipulating the CS–US interval and the US intensity on the acquisition of the classically conditioned nictitating membrane response. The CS was a 50 ms tone and the US was a 50 ms electric shock. The ISI values were 125, 250, 500, and 1000 ms. The fact that conditioning occurred at ISIs much larger than CS duration implies that an internal trace of the CS, which we have called I_{CS} , is stored in short-term memory subsequent to CS offset, as in Figure 2. The US intensities were 1, 2, and 4 mA.

Smith (1968) 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. In addition, an increase in the mean of the peak response time correlated with an increase in the variance of the response curve, for each ISI.

All of these properties are evident in the computer simulation of Figure 7. The absolute sizes of the empirically measured responses increase slower-than-linearly in Figure 8 as a function of shock intensity, rather than linearly as in the computer simulations in Figure 7. This fact suggests that shock intensity is transformed by a slower-than-linear signal function in vivo, rather than the linear signal function that we used to engage the activation spectrum of the model. Such a slower-than-linear transformation can easily be generated by a preprocessing step at which the CS is averaged by a shunting on-center off-surround feedback network at which the CS is stored in short-term memory (Grossberg, 1982, 1988). The output from this short-term memory representation to the timing circuit is I_{CS} .

The qualitative similarities between the data in the top panel of Figure 8 and the computer simulation in Figure 7a are quantified in Figures 9 and 10. Figure 9 plots data points and computer simulations together. Figure 10 plots four measures of data and simulation at ISI values of 0, 125, 250, 500, and 1000




FIGURE 5. Generation of the response function $R(t)$. The CS was presented at the beginning of each learning trial. The US was presented 400 ms later (thus the ISI = 400) and kept on for 50 ms during 4 learning trials, which were followed by one test trial during which only the CS was presented. (a) Graph of the output signal $R(t)$ through time on all five trials. (b) After each trial, the time scale was reset to $t = 0$ to superimpose the output signal with a common initial time. The sudden jump in four of the five curves is due to the I_{US} . All the output curves are centered at the ISI because the output threshold $F = 0$ in (5). If F is chosen positive, the successive output curves move progressively backwards in time and become progressively better centered at the ISI as learning proceeds. (c) All the doubly gated signals $h_i(t) = f_i(x_i(t))y_i(t)z_i(t)$, $i = 1, 2, \dots, 80$, are plotted through time on the fifth trial. (d) All the gated signals $g_i(t) = f_i(x_i(t))y_i(t)$, $i = 1, 2, \dots, 80$, are plotted through time on the fifth trial. Parameters are chosen as in Figure 4.




FIGURE 6. As in Figure 5b, superimposed plots of the output signal $R(t)$ on four successive learning trials and one performance trial are shown, along with plots of all the doubly gated signals $h_l(t)$, $l = 1, 2, \dots, 80$, on the fifth trial. Each panel displays the results at a different ISI: (a) $ISI = 0$ ms; (b) $ISI = 500$ ms; and (c) $ISI = 1000$ ms.

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





FIGURE 7. Inverted U in learning as a function of ISI. The output signal functions $R(t)$ are plotted on a test trial, in response to the CS alone, subsequent to 10 prior learning trials with CS-US separated by different ISI's. Successive curves from left to right were generated by ISI's of 0 (the lowest amplitude curve), 125, 250, 500, and 1000 ms using a US duration of 50 ms. Two different I_{us} intensities were used in (a) and (b), respectively. In (a), $I_{us} = 10$. In (b), $I_{us} = 5$. All other parameters were chosen as in previous figures.

ms. The four measures are peak time (μ), standard deviation (σ), Weber fraction (W), and peak amplitude (A). Peak time (μ) was defined as the time at which the response amplitude reached its maximum value at each ISI. Standard deviation (σ) was estimated by approximating each response curve by a normal distribution and determining the times at which the amplitude was equal to .61 of the curve's peak value. This criterion was chosen because the

interval between the times at which response amplitude equals .61 of its peak value is approximately 2σ in length. To see this, consider a normal distribution $1/\sqrt{2\pi}\sigma \exp[-(t - \mu)^2/2\sigma^2]$. Its amplitude when $|t - \mu| = \sigma$ is $1/\sqrt{2\pi}\sigma \exp(-1/2)$. Its amplitude when $t = \mu$ is $1/\sqrt{2\pi}\sigma$. The ratio of these amplitudes is $\exp(-1/2) \cong .61$. The Weber fraction W was defined as $W = \sigma/\mu$.

Despite the coarse nature of these approxima-




FIGURE 8. Conditioning data from a nictitating membrane learning paradigm. Mean topography of nictitating membrane response after learning trial 10 with a 50 ms CS, ISI's of 125, 250, 500, and 1000 ms, and different (1, 2, 4 MAmp) intensities of the shock US in each subsequent panel. Reprinted from Smith (1968) with permission.

tions, Figure 10 reveals a remarkably good fit between experimental and simulated values of all the parameters μ , σ , W , and A at all the reported ISIs. Of particular interest is the approximately constant value of the Weber fraction W as a function of ISI, in particular its tendency to approach a positive asymptote with increasing values of the ISI (Killeen & Weiss, 1987).

Although the Spectral Timing Model provides a good quantitative fit to conditioning data acquired over a relatively small number of trials, say 1–20, the associative learning equation (4) needs to be made slightly more complex to work well over very large numbers of trials. This is true because all z_i for which $f(x_i)y_i > 0$ during times when $I_{US} > 0$ can approach I_{US} , albeit at different rates, as $t \rightarrow \infty$. Adding a very slow passive decay term $-\epsilon z_i$ to eqn (4) can overcome this potential difficulty.

7. INVERTED U IN LEARNING AS A FUNCTION OF ISI

A basic property of both the simulated response functions $R(t)$ in Figure 7 and the data summarized in Figure 8 is an inverted U in learning as a function of the ISI. In other words, there exists a positive ISI that is optimal for learning. In Figure 7, this optimal ISI is approximately 250 ms. Learning is weaker at both smaller and larger values of the ISI.




FIGURE 9. Comparison of experimental data from Figure 8 (US intensity equal 1 MAmp) with computer simulation from Figure 7a (I_{US} intensity equal 10). Simulated values are computed as curves interpolated through 10 values around the time of US presentation separated by $.16 \times \text{ISI}$. Open circles represent experimental data computed at the same times corrected to make the peak values of the experimental and simulated 250 ms ISI curves coincide. The correlation between simulated and experimental points equals $r = .835$.




FIGURE 10. Comparison between experimental and simulated peak time (μ), standard deviation (σ), Weber fraction (W), and peak amplitude (A). See text for details. The correlation between simulated and experimental points for μ is $r = .9996$ ($p < .001$), for σ is $r = .9761$ ($p < .005$), and for A is $r = .9666$ ($p < .01$).

A number of experimental conditions have been developed to better understand this fundamental property. In simultaneous conditioning (zero ISI), CS and US begin together. In delay conditioning, the CS precedes the US, and the US overlaps the CS. In trace conditioning, the CS precedes the US, and the US is presented after the CS offset. Conditioning is typically more efficacious when the CS precedes the US than when the two are presented together (Gormenzano, Kehoe, & Marshall, 1983).

It has been found that different response systems in a given species present different optimal ISIs. As illustrated above, the nictitating membrane conditioned response in rabbits has an optimal ISI of around 250 ms (Smith, 1968). Heart rate conditioning in rabbits is optimal with a 7-s ISI (Schneiderman, 1972). Conditioned leg flexion in cats is optimal with a 500-ms ISI (McAdam, Knott, & Chiorini, 1965). Salivary conditioning in dogs is optimal with a 20-s ISI (Konorski, 1948). Conditioned licking in rats is optimal with a 3-s ISI (Boice & Denny, 1965). Heart rate conditioning in rats is optimal with a 5-s ISI (Black & Black, 1967).

Although the Spectral Timing Model successfully generates such a positive optimal ISI, it seems clear that this circuit is not the only one subserving the optimal ISI that is behaviorally observed.

This can be seen by considering the phenomenon

of secondary excitatory conditioning. In secondary excitatory conditioning, two CSs are employed; call them CS_1 and CS_2 . Let CS_1 be conditioned with a US until CS_1 can elicit some of the reinforcing properties of the US. Then present the two CSs simultaneously as a compound stimulus $CS_1 + CS_2$. The conditioning of CS_2 to the new reinforcer CS_1 is much attenuated relative to the conditioning that would have occurred if CS_1 was presented before CS_2 .

On the other hand, consider an experiment in which CS_1 and CS_2 are equally salient to the organism and the compound cue $CS_1 + CS_2$ is presented before a US on conditioning trials. Then both CS_1 and CS_2 can be effectively conditioned to the US.

Thus the attenuation in the conditioning of CS_2 to CS_1 when CS_1 and CS_2 are simultaneously presented and CS_2 has previously been conditioned to US cannot be due merely to the *simultaneity* of CS_1 and CS_2 in their capacity as sensory events. Rather it must be due to the effects of reading-out within the network the *reinforcing properties* of CS_2 by the sensory representation of CS_2 .

A model capable of explaining how such attentional blocking of CS_2 by a simultaneous conditioned reinforcer CS_1 is outlined in Part II. Computer simulations of attentional blocking within this model are found in Grossberg and Levine (1987; reprinted in Grossberg, 1988).

8. MULTIPLE TIMING PEAKS

Another functionally useful model property that matches experimental conditioning data concerns the ability of a single CS to read out responses at a series of learned delays. This multiple timing property provides strong indirect evidence that each CS sends signals to a complete activation spectrum, rather than to a single tunable delay.

Figure 11 depicts the outcome of a computer simulation in which a CS is paired with a US whose ISI is chosen on alternate trials at two different values. When the CS is subsequently activated on a recall trial, the response function $R(t)$ generates two peaks, with each peak centered at one of the ISIs.

The parameters used in the simulation of Figure 11 are the same as those used to fit the data in Figure 10 concerning response time, amplitude, standard deviation, and Weber fraction. It is therefore of particular interest that the model simulations in Figure 11 strikingly resemble the multiple timing data of Millenson, Kehoe, and Gormenzano (1977) that are summarized in Figure 12.

Millenson, Kehoe, and Gormenzano (1977) presented rabbits in a nictitating membrane paradigm with a tone CS followed by a shock US at two randomly alternating ISIs of 200 and 700 ms. The CS terminated at US onset, and the US had a 50 ms duration. Each row in Figure 12 corresponds to a different experimental condition. The experiment

summarized in row 1 used a 200 ms ISI throughout. The experiment in row 5 used a 700 ms ISI throughout. Compare these relative peak times, amplitudes, and Weber fractions with the model simulation in Figure 11.

Experiments summarized in the middle three rows used varying fractions of the two ISI delays during learning trials. In the second row, the ISI equaled 200 ms on 7/8 of the learning trials and 700 ms on 1/8 of the learning trials. In the third row, the ISI equaled each of these values on 1/2 of the learning trials. In row four, the ISI equaled 200 ms on 1/8 of the trials and 700 ms on 7/8 of the trials.

Each column in Figure 12 corresponds to a different test condition subsequent to a set of learning trials. During such a test, a CS, but no US, was presented. In column 1, the CS duration was 200 ms. In column 2, the CS duration was 700 ms. In each panel, a test curve is displayed after 3 days and after 10 days of prior learning.

The data curves of greatest interest are in row 3, column 2. These curves are strikingly similar to the model simulation in Figure 11. Row 3, column 1 is also of interest, because it shows that termination of a CS of 200 ms duration under these conditions prevents strong perseveration of its I_{CS} curve for the additional 500 ms needed to read out a large response at 700 ms.

The parameters used to fit the data in Figures 9, 10, and 12 generate broadly tuned timing peaks.

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




FIGURE 11. Multiple timing peaks due to learning with more than one ISI. The output signal function $R(t)$ is plotted on a test trial after 20 learning trials during which a US of intensity 10 was presented alternately at an ISI of 200 ms and 800 ms.




FIGURE 12. Conditioning data from the nictitating membrane learning paradigm of Millenson, Kehoe, and Gormenzano (1977). Data shown after learning trials 3 and 10 using a tone CS of duration 200 msec and 700 msec, ISI's of 200 msec and 700 msec, and a shock US of 50 msec duration. See text for details. Reprinted with permission.

Finer peaks can, however, be generated, should technological applications so require, without disturbing other useful qualitative properties. For example, using a different set of parameters, the simulation reported in Figure 13 generates the same qualitative series of peaks as in Figure 7; but a sharper multiple timing curve (Figure 13b) than in Figure 11.

9. EFFECT OF INCREASING US DURATION

Figure 14 depicts the results of a simulation that illustrates the effects of increasing US duration upon the response $R(t)$. The I_{US} intensity was twice as large in Figure 14a than in Figure 14b. A zero ISI was employed throughout. Two effects are generated: a shift of peak time to a value towards the midpoint of the US, and an overall increase in conditioned response.

Burkhardt and Ayres (1978) have collected analogous data (Figure 15) in a paradigm wherein rats were presented with an auditory CS and a simultaneous (zero ISI) shock US. When the CS was later presented while the rats were drinking water, the CS presentation elicited a suppression of licking whose relative magnitude before and after CS onset (the suppression ratio) was used to measure the strength of the conditioned fear elicited by the CS. During conditioning, a grid-shock US of 2, 4, or 8 s duration began simultaneously with a noise CS of 2, 4, or 8 s duration in the combinations 2-2, 2-4, 4-4, 4-8, and 8-8. As in Figure 14, Burkhardt and Ayres (1978) found that conditioning increased as a function of




FIGURE 13. Multiple timing peaks due to learning with more than one ISI. The parameters were chosen as in Figure 4, with the exception of $\beta = .4$.




FIGURE 14. Effect of increasing US duration and intensity on learning. In both (a) and (b), ISI = 0. The output signal $R(t)$ is plotted on a test trial after 4 learning trials using I_{US} functions of duration 50, 100, and 150 ms. In (a), US intensity equals 10. In (b), US intensity equals 5.

US duration, as well as of CS-US overlap, another property easily explained by the model.

10. EFFECT OF INCREASING CS INTENSITY

Figure 7 showed that an increase of US intensity alters the amplitude of the response functions $R(t)$, but not their timing, and Figure 8 showed that the conditioning data of Smith (1968) conform to these properties. A different pattern of results is obtained if the CS intensity is altered. Figure 16 illustrates a computer simulation in which the system was trained with a CS and US of constant intensity and an ISI of 800 ms across learning trials. The Figure shows that a test trial using the same CS intensity generates a response function $R(t)$ that peaks at 800 ms, but

a test trial using a CS of twice that intensity generates a response function $R(t)$ that peaks at 400 ms. Thus, increasing CS intensity "speeds up the clock" that calibrates the response reaction time. Such a speed-up is a straightforward consequence of eqn (1).

Section 17 describes experimental data which are consistent with these properties. In order to analyze these data, we first need to explain how the timing circuit is embedded within a larger architecture that controls the stable self-organization of cognitive-emotional representations.

PART II TIMED GATING OF READ-OUT FROM THE ORIENTING SUBSYSTEM

11. LOCATING THE TIMING CIRCUIT WITHIN A SELF-ORGANIZING SENSORY- COGNITIVE AND COGNITIVE- REINFORCEMENT ART NETWORK

The timing circuit is hypothesized to form part of interacting sensory-cognitive and cognitive-reinforcement circuits which have been progressively developed since the late 1960s to explain behavioral and neural data about recognition, reinforcement, and recall.

Sensory-cognitive interactions in the theory are carried out by an Adaptive Resonance Theory (ART) circuit (Carpenter & Grossberg, 1987a, 1987b, 1988; Cohen & Grossberg, 1986, 1987; Grossberg, 1976, 1982, 1987; Grossberg & Stone, 1986). Such ART architectures are designed to explain how internal representations of sensory events, including conditioned stimuli (CS) and unconditioned stimuli (US), are learned in real-time in a stable fashion in response to noisy, nonstationary environments.

As in Figure 17, a sensory-cognitive ART circuit is broken up into an attentional subsystem and an orienting subsystem. The attentional subsystem learns ever more precise internal representations of and responses to events as they become more familiar. The attentional subsystem also learns the top-down expectations that help to stabilize memory of the learned bottom-up codes of familiar events. The orienting subsystem resets the internal representation that is active in short-term memory (STM) in the attentional subsystem when an unfamiliar or unexpected event occurs. The orienting subsystem also energizes the orienting response, including the movements triggered by novel events that enable such events to be more efficiently processed.

The orienting subsystem is activated when a sufficiently large mismatch occurs within the attentional subsystem between bottom-up sensory input signals and learned top-down expectations. In Figure 17a, the learned top-down expectations are read-out from

MEDIAN CS TIMES IN SECONDS
FOR EACH GROUP IN EXPERIMENTS 1, 2, AND 3

Group	Baseline-Lick Last Training	Test 1	Test 2
Experiment 1			
0	1.6	5.2	3.2
1	1.6	3.4	3.4
4	1.9	37.8	6.2
64	1.4	2.7	4.8
128	1.6	8.8	7.8
Experiment 2			
0	1.8	3.2	1.9
1	1.7	15.1	4.8
2	1.9	12.9	3.8
4	1.7	94.4	6.4
8	1.6	64.5	8.3
Experiment 3			
2-2	1.5	4.9	4.7
2-4	1.5	15.6	8.9
4-4	1.4	30.6	5.4
4-8	1.5	58.2	9.9
8-8	1.4	107.8	7.0

FIGURE 15. Data of Burkhardt and Ayres (1978) on conditioning an auditory CS and a simultaneous (zero ISI) shock US. See the text for details. Reprinted with permission.

level F_2 to level F_1 , and matching of expectations with bottom-up input patterns occurs at level F_1 . When a mismatch occurs, the orienting subsystem A is activated and causes an STM reset wave to be delivered to level F_2 . This STM reset wave resets the sensory representations of cues that are currently being stored in STM at F_2 .

As noted in Section 3, one function of the timing circuit is to prevent spurious resets of active internal

representations in response to mismatches due to *expected nonoccurrences* of sensory events. In addition, the timing circuit should not prevent registration of bottom-up input patterns and their matching with active top-down expectations. Thus the timing circuit does not interfere with processing within the attentional subsystem.

Instead, we hypothesize that the timing circuit inhibits read-out of the STM reset wave from the ori-




FIGURE 16. Effect of CS intensity on "clock speed." After 10 learning trials were carried out with an ISI = 800 msec, 2 test trials were carried out, one with the CS intensity of 1 used in training, whose output $R(t)$ peaked at 800 msec, and one with a CS intensity of 2 that caused the output $R(t)$ to peak at 400 msec.




FIGURE 17. Anatomy of an adaptive resonance theory (ART) circuit: (a) Interactions between the attentional and orienting subsystems: Learning of recognition codes takes place at the long term memory (LTM) traces within the bottom-up and top-down pathways between levels F_1 and F_2 . The top-down pathways can read-out learned expectations, or templates, that are matched against bottom-up input patterns at F_1 . Mismatches activate the orienting subsystem A , thereby resetting short term memory (STM) at F_2 and initiating search for another recognition code. Output from subsystem A can also trigger an orienting response. Sensitivity to mismatch at F_1 is modulated by vigilance signals from the drive representations. (b) Trainable pathways exist between level F_2 and the drive representations. Learning from F_2 to a drive representation endows a recognition category with conditioned reinforcer properties. Learning from a drive representation to F_2 associates the drive representation with a set of motivationally compatible categories.

enting subsystem A (Figure 17a). Thus when the timing circuit is active, both STM reset within the attentional subsystem and the orienting response are inhibited. When the timing circuit is inactive, an *unexpected nonoccurrence* of an event is able to trigger the STM reset and orienting response needed to cope with the unexpected event.

To analyze how the timing circuit works, we summarize some basic properties of another part of the attentional subsystem. This is the network which controls the learned interactions between recognition and reinforcement mechanisms that focus attention upon motivationally salient events. We assume, in particular, that the timing circuit forms part of the interaction from cognitive to reinforcement repre-

sentations whereby sensory cues learn to become conditioned reinforcers (Figure 17b).

12. COGNITIVE-REINFORCEMENT CIRCUIT

Recognition is only one of several processes whereby an intelligent system can learn a correct solution to a problem. Reinforcement and recall are no less important in designing an autonomous intelligent system.

Reinforcement, notably reward and punishment, provides additional information in the form of environmental feedback based on the success or failure of actions triggered by a recognition event. Reward and punishment calibrate whether the action has or has not satisfied internal needs, which in the biological case include hunger, thirst, sex, and pain reduction, but may in machine applications include a wide variety of internal cost functions. Reinforcement can modify the formation of recognition codes and can shift attention to focus upon those codes whose activation promises to satisfy internal needs based upon past experience. For example, both green and yellow bananas may be recognized as part of a single recognition category until reinforcement signals, contingent upon eating these bananas, differentiates them into separate categories.

Recall can generate equivalent responses or actions to input events that are classified by different recognition codes. For example, printed and script letters may generate distinct recognition codes, yet can also elicit identical learned naming responses.

The type of ART circuit depicted in Figure 17a is devoted entirely to the stable self-organization of sensory and cognitive recognition codes. Feedback interactions among recognition and reinforcement circuits, as in Figure 17b, are also posited by the theory, and in fact were the first type of ART circuit to be defined (Grossberg, 1975, 1982). In these applications, the circuit at which recognition codes are processed is called a *sensory representation S* , and the circuit at which reinforcement and homeostatic, or drive, signals are processed is called a *drive representation D* (Grossberg, 1971, 1972b, 1987), as in Figure 17b. Thus a reinforcing event, such as a reward or punishment, possesses both a sensory representation in its capacity as a sensory event, and a drive representation in its capacity as a motivationally significant reinforcer.

During classical conditioning, a familiar conditioned stimulus (CS) may initially have a sensory representation S , but no drive representation D . Pairing a CS with an unconditioned stimulus (US) that does have reinforcing properties causes several types of learning to occur. In particular, repeated pairing of a CS sensory representation, S_{CS} , with activation of a drive representation, D , by a US rein-

forcer causes the modifiable synapses connecting S_{CS} with D to become strengthened. This conditioning process converts the CS into a *conditioned reinforcer* (Figures 17b and 18). Incentive motivation pathways from the drive representations are also assumed to be conditionable. These conditioned $S \rightarrow D \rightarrow S$ feedback pathways shift attention to focus upon the subset of active sensory representations which have been previously reinforced and are motivationally compatible. This shift of attention occurs because the sensory representations which emit conditioned reinforcer signals $S \rightarrow D$ and receive conditioned incentive motivation signals $D \rightarrow S$ compete among themselves for a limited capacity short-term memory (STM) via on-center off-surround interactions (Figure 18). When incentive motivational feedback signals are received at the sensory representational field, these signals can bias the competition for STM activity towards motivationally salient cues. More generally, such feedback interactions between S and D can reorganize the STM pattern across S to be compatible with reinforcement constraints. This STM pattern can then be incorporated through learning into the sensory-cognitive recognition code via an ART circuit of the type shown in Figure 17.

In order to explain the moment-by-moment dynamics of conditioning, an additional microcircuit needs to be embedded in the drive representations of the macrocircuit depicted in Figure 18. This microcircuit, called a *gated dipole* (Grossberg, 1972a,




FIGURE 18. Schematic conditioning circuit: Conditioned stimuli (CS_i) activate sensory representations (S_{CS_i}) which compete among themselves for limited capacity short-term memory 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.

1972b), instantiates a neurophysiological theory of opponent processing. The need for a certain type of opponent processing for conditioning circuits can be seen from the following considerations.

13. THE GATED DIPOLE OPPONENT PROCESS

In the cognitive-reinforcement circuit, CS 's can become conditioned reinforcers by being associated with either the onset or the offset of a reinforcer. For example, a CS that is conditioned to the onset of a shock can become a source of conditioned fear (excitor). A CS that is conditioned to the offset of a shock can become a source of conditioned relief (inhibitor). A gated dipole opponent process explains how the offset of a reinforcer can generate an off-response, or antagonistic rebound, to which a simultaneous CS can be conditioned. A gated dipole is a minimal neural network opponent process which is capable of generating a sustained, but habituated, on-response (e.g., a fear reaction) to onset of a cue (e.g., a shock), as well as a transient off-response (e.g., a relief reaction), or antagonistic rebound, to offset of the cue. The on-responses are processed through the on-channel D^+ of the gated dipole, whereas the off-responses are processed through the off-channel D^- of the gated dipole. In addition, such a gated dipole must be joined to a mechanism of associative learning, whereby CS 's learn to become conditioned excitors via $S \rightarrow D^+$ learning and conditioned inhibitors via $S \rightarrow D^-$ learning.

14. ADAPTIVE TIMING AS SPECTRAL CONDITIONED REINFORCER LEARNING

The feedforward adaptive timing circuit is assumed to be a variant of $S \rightarrow D^+$ conditioned reinforcer learning. The main new idea is that the on-channel's population of neurons D^+ is broken up into neuron subpopulations whose membrane properties enable them to respond to inputs at different rates α_i , as in eqn (1). In other words, by selecting a sloppy parametric specification of cell reaction rates, nature can discover an adaptive timing mechanism—if such sloppiness is permitted at the proper processing stage of a gated dipole circuit!

Once this is achieved, standard gated dipole mechanisms respond to the activation spectrum α_i in the manner described in Part I. In particular, eqn (2) for the transmitter gate is the standard gating equation that gave a gated dipole its name; eqn (3) is a variant of the dipole's standard associative learning law; and eqn (4) simply computes the total output from all subpopulations of the dipole's on-channel. We sum-

marize this fact by saying that adaptive timing is a type of *spectral conditioned reinforcer learning*.

15. TIMED INHIBITION OF THE ORIENTING SUBSYSTEM BY DRIVE REPRESENTATIONS

It remains to explain how a timing circuit embedded within the on-channels of gated dipole drive representations achieves the functional properties described in Section 3. These properties follow if we assume, in addition, that the drive representations D inhibit the orienting subsystem A , as in Figure 19. In Figure 17a, level F_1 also inhibits A . Thus several processing levels within the attentional subsystem are assumed to inhibit the orienting subsystem. The hypothesis of competition from D to A representations was first made within the context of ART-type models in Grossberg (1975; reprinted in Grossberg, 1982, pp. 284–286). Given $D \rightarrow A$ inhibition, spectral conditioned reinforcer learning generates the desired adaptive timing properties as follows.

After CS-US conditioning at a fixed ISI, presentation of the CS activates its sensory representation S_{CS} , which activates its conditioned drive representation D with a response curve $R(t)$ of the form depicted in Figure 7. Each of these response curves $R(t)$ begins to grow right after I_{CS} read-out; remains positive throughout an interval whose total width covaries with the ISI, due to the approximate constancy of the Weber fraction W (Figure 10); and peaks at the ISI. Inhibition of A by D thus prevents STM reset by expected nonoccurrences of the US throughout a time interval that is centered at the




FIGURE 19. Inhibition of the orienting subsystem A by the output from a drive representation D . The Spectral Timing Model is assumed to be part of the network whereby conditioning of a sensory representation S to a drive representation D endows S with conditioned reinforcer properties. As S reads-out spectrally timed conditioned signals to D , D inhibits output signals from A and thereby prevents expected nonoccurrences of the US from resetting STM and triggering orienting responses.

expected delay of the US whose width covaries with this delay.

16. TIMED ACTIVATION OF THE HIPPOCAMPUS AND THE CONTINGENT NEGATIVE VARIATION

Because it is activated by the drive representations D , positive feedback from D to S along the $D \rightarrow S$ incentive motivational pathways is also timed to provide peak motivational support for release of a conditioned response (Figure 18) at the expected delay of the US.

In Grossberg (1975, Section VII, and 1978, Section 16; reprinted in 1982), such $D \rightarrow S$ feedback was first interpreted to be a formal analog of the contingent negative variation, or CNV, event-related potential. The CNV had earlier been experimentally shown to be sensitive to an animal's expectancy, decision (Walter, 1964), motivation (Cant & Bickford, 1967; Irwin, Rebert, McAdam, & Knott, 1966), preparatory set (Low, Borda, Frost, & Kellaway, 1966), and arousal (McAdam, 1969). It is also a conditionable wave whose timing tends to match the ISI. Until the present work, development of our conditioning theory, as summarized in Grossberg (1987, Chapter 1, Sections 23 and 25, and Chapter 2, Sections 30, 43, 53, 57, and 60), suggested how the CNV is conditioned and how it is related to expectancy, decision, motivation, preparatory set, and arousal. The theory had not, however, heretofore explained how the learning process enables CNV timing to mimic the ISI. The present extension of the theory provides an explanation through the hypothesis of spectral conditioned reinforcer learning. The interpretation of drive representations in terms of hypothalamo-hippocampal interactions (Grossberg, 1971, 1982, 1987) provides an anatomical marker for directly testing the existence of spectral activation.

The hypotheses that drive representations include hippocampus and that the hippocampus is involved in conditioned timing have also received support from neurophysiological experiments (Berger & Thompson, 1978; Delacour & Houcine, 1980; Hoechler & Thompson, 1980; Rawlins, 1985; Rawlins, Feldon, & Gray, 1982; Solomon, 1979, 1980; Solomon, van der Schaff, Thompson, & Weisz, 1986).

17. EFFECT OF CS INTENSITY ON TIMED MOTOR BEHAVIOR

We are now ready to use the model property illustrated in Figure 16 to suggest how changes in CS intensity and various drug manipulations may cause observed changes in certain timed motor behaviors of animals.

For example, changes in CS intensity alter the

conditioned key pecking behavior of pigeons (Wilkie, 1987). In these experiments, each pigeon was pretrained to discriminate between short (2 s) and long (10 s) houselight presentations. In one set of 30 sessions, a bright houselight was used. In another, a dim houselight was used. In all learning sessions, the 20-s intertrial interval was spent in complete darkness. In each of the 80 trials in a session, the probability was .5 that the houselight presentation was short.

Immediately after the short or long houselight presentation was completed, two pecking keys were lit, one with red light and the other with green light. The right-left locations of the red and green keys was varied randomly over trials. For some pigeons, red was designated as the correct key to peck after a short stimulus, and green was the correct key to peck after a long stimulus. For other pigeons, the colors were reversed. Pecking of the correct key produced 5-s access to mixed grain on a partial 75% reinforcement schedule.

During the experiment proper, all pigeons received approximately 35 sessions, each comprising 80 trials. On one-half of the trials (randomly determined), the houselight was bright. On the other trials, it was dim. On a quarter of all trials (randomly determined), the light presentation was 2 s in duration. On another quarter, it was 10 s. On the remaining trials, it was equally probable that the light would be 4, 6, or 8 s in duration. Thus there were 10 types of trials in total: 2, 4, 6, 8, and 10 s bright lights, and 2, 4, 6, 8, and 10 s dim lights, presented in randomized order. Correct choices on 2 and 10 s trials of both dim and bright lights always produced 5-s access to grain. Choices on 4, 6, and 8 s trials were never rewarded.

In each session, a record was kept of the number of times the "short" choice key was selected when 2, 4, 6, 8, and 10 s lights were presented. These values were accumulated over sessions and used to calculate the percentage of trials on which pigeons chose the "short" alternative after durations of 2, 4, 6, 8, and 10 s lights. These measures were calculated separately for bright and dim light trials.

It was found that the pigeons chose the "short" key more frequently in response to longer durations of the dim light. In other words, a dim light slowed down the time scale, as in Figure 16. Wilkie (1987, p. 38) noted that "it is not intuitively obvious how intensity would affect something like a counter or how any such effect would be manifested in dim signals' being perceived as being shorter."

In order to provide a more detailed explanation of these data based upon the model property illustrated in Figure 16, several properties of the experiment need to be kept in mind. In particular, the presentation of the red and green keys immediately

followed the short or long houselight stimulus, and reward or non-reward immediately followed a correct key peck. Suppose that an internal representation of a dim or bright houselight activated a full activation spectrum, and that food reward caused conditioning of those spectral populations that were active when the reward occurred, as in Figure 4. In response to a short CS, only rapidly reacting spectral populations could become conditioned. In response to a long stimulus, only those spectral populations which became active after a longer CS duration could become conditioned. Thus the basic properties of the timing model explain how, in response to a CS of any fixed intensity, only a properly timed subset of spectral populations could become conditioned. In addition, the model property depicted in Figure 16 shows how dimming of the CS can, other things being equal, slow down the read-out of the clock.

Further discussion is needed, however, to explain how dim and bright houselights are discriminated in the first place, and how differential reward of both dim and bright short lights and of dim and bright long lights generated the main effect that pigeons peck the "short" key in response to longer durations of dim light. Indeed, *both* the dim light and the bright light are conditioned to different key pecks based on their *duration*, not their *intensity*. Why should longer dim lights tend to generate the key peck that was associated with a short duration light independent of its intensity?

The computer simulations reported in Figure 16 would imply this result if some of the spectral cells that are activated by a *short bright* light are also activated by a *long dim* light. On those learning trials when these cells are activated by a short bright light, they would be conditioned, via incentive motivational feedback signals (Figure 18), to the internal representation of the key that signifies a short stimulus. On those learning trials when these cells are activated by a long dim light, they would amplify these internal representations and thereby favor this key in the STM competition for which key the pigeon will attend and thereupon peck.

18. SPATIAL CODING OF STIMULUS INTENSITY BY A PTS SHIFT MAP

We trace this property to the manner in which different intensities of the same stimulus are discriminated by the animal. Suppose that a particular stimulus input, such as a white light, is coded by a population of cells. Grossberg and Kuperstein (1986, pp. 160-167) have developed a model of such a coding population in which different input intensities maximally activate different subsets of the total population. Thus, input *intensity* is recoded into the maximally activated *spatial location* within the pop-

ulation. Since distinct subsets of the population can activate different output pathways, different input intensities can control their own spectral populations, and can be conditioned to activate the drive representations at different times.

The Grossberg and Kuperstein (1986) model is called a Position-Threshold-Slope (PTS) Shift Map. To generate this spatial map of input intensity, the cells within the population are assumed to possess different output thresholds and different sensitivities to input increments. Cells with higher thresholds are assumed to be more sensitive. Thus, essentially all input intensities (e.g., dim and bright lights) can generate output signals from cells with low output thresholds, whereas only high input intensities (e.g., bright lights) can generate output signals from cells with high thresholds. Due to the greater sensitivity of the high-threshold cells, the spatial locus of maximal activation changes with input intensity.

Populations of cells whose output thresholds and input sensitivities covary have been found in the abducens and oculomotor nuclei (Luschei & Fuchs, 1972; Robinson, 1970; Schiller, 1970). The present analysis suggests that such populations may also exist in thalamocortical sensory processing areas.

19. EFFECT OF DRUGS ON TIMED MOTOR BEHAVIOR

Wilkie (1987, p. 38) has speculated, based on earlier results of Maricq, Roberts, and Church (1981), that "drug and light-intensity effects might both be mediated by a state of arousal that affects the pacemaker rate." Meck and Church (1987) have reviewed a number of experiments, including the Maricq et al. (1981) and Meck (1983) experiments, and have collected additional data on the effects of drugs on timed motor behavior. The major properties of these drug manipulations are consistent with the Spectral Timing Model.

Meck and Church (1987) noted that an increase in the effective level of brain dopamine at the synapse increases clock speed and that a decrease in the effective levels of brain dopamine decreases clock speed. Methamphetamine and L-dopa increase dopamine at the synapse and change timing functions in a manner that can be interpreted as an increase in clock speed. Neuroleptics, such as haloperidol, which block dopamine receptors change timing functions in a manner that can be interpreted as a decrease in clock speed.

Likewise, experimental evidence suggests that an increase in the effective level of brain acetylcholine at the synapse reduces the remembered time of reinforcement in long-term memory, and thus speeds up the clock in short-term memory. A decrease in the effective level of brain reinforcement increases

the remembered time of reinforcement in long-term memory, and thus slows down the clock in short-term memory. For example, both physostigmine and phosphatidylcholine change timing in a manner interpretable as a decrease in remembered time of reinforcement, whereas atropine and aging cause an increase.

Since the introduction of gated dipole theory in 1972 (Grossberg, 1972b; reprinted in Grossberg, 1982), it has been predicted that the habituating transmitter gates, as defined in eqn (2), are chemically realized in the brain by a catecholamine, such as dopamine or norepinephrine, and that the long-term memory traces, as defined in eqn (4), are realized in the brain by acetylcholine. Thus the present model is consistent with the recent dopamine and acetylcholine data if a gated dipole circuit exists that processes the CS input before it generates I_{CS} in eqn (1). In this way, the aforementioned drug manipulations would alter the intensity of the CS, and thereby speed up or slow down the clock in the manner indicated in Figure 16.

Such habituating and LTM transmitter systems are, in fact, postulated as part of the adaptive coding circuitry that self-organizes an internal representation of the CS in an ART circuit such as that depicted in Figure 17 (Grossberg, 1982, 1987).

These drug data also raise the question whether the habituating transmitter and the LTM transmitter *within* the Spectral Timing Model itself can influence clock speed. This would be the case if the timing circuit included internal feedback loops whereby the two types of transmitters feed back their influence to the spectral activities defined in eqn (1). Such feedback pathways have previously been postulated to exist in the gated dipole circuits that regulate the learning of conditioned reinforcers (Grossberg, 1982; Grossberg & Schmajuk, 1987), of which the Spectral Timing Model is herein assumed to be a specialization. It remains for future research to determine how a Spectral Timing Model with internal feedback pathways may be designed.

20. CONCLUDING REMARKS: TIMING PARADOX AND MULTIPLE TYPES OF TIMING CIRCUITS

There exist multiple types of timing mechanisms in the brain. The present article considers only the type of timing that enables an organism to time and differentially respond to an expected nonoccurrence, an expected occurrence, and an unexpected nonoccurrence of a sensory event subsequent to a prior sensory event or action.

In so doing, the article clarifies a Timing Paradox that becomes apparent upon closer inspection of this type of timing problem. On the one hand, in response

to any fixed choice of conditionable ISI, it is desired that the learned *optimal* response delay approximate the ISI. Thus the model must be capable of an accurate discrimination of individual temporal delays. On the other hand, it is also desired that spurious orienting responses be inhibited in response to expected nonoccurrences that may occur *throughout* the ISI interval subsequent to a CS onset. Thus the inhibitory signal must be temporally distributed throughout the ISI interval.

The Spectral Timing Model reconciles the two requirements of accurate optimal temporal delay and temporally distributed activation via the Weber law property (Section 6). According to this property, the standard deviation of the model response scales with its peak time. Consequently, the model begins to immediately generate an output signal that may be used to inhibit the orienting subsystem, even though its peak output is accurately located at the ISI.

This key property distinguishes the Spectral Timing Model from a model that uses conditionable pathways with brief sampling signals and variable delay lines to learn to time the ISI delay. In such a model, use of a single ISI during training would lead to a zero learned output in response to the CS until the ISI had elapsed. The output from such a model could not be used to inhibit orienting responses in response to expected nonoccurrences.

The Spectral Timing Model is also not mechanistically the same as model circuits which have been identified to self-organize the learning and long-term memory of serially ordered behaviors, or the encoding of event sequences in short-term memory, or the encoding of sequential rhythmic properties in short-term memory, or the clock-like oscillatory timing of circadian rhythms. In particular, the type of timing controlled by the Spectral Timing Model occurs within hundreds of milliseconds or a few seconds at most of a single behavioral response. It is not the type of timing that may be spread over many seconds or minutes whereby sequences of behavioral acts are regulated. Neural network models for these alternative timing capabilities have been described in the books Grossberg (1982, 1987) and Grossberg and Kuperstein (1986).

For example, the Spectral Timing Model, at least in its present form, cannot explain how an animal can learn to interrupt a timed behavioral sequence during a signalled time-out period and continue the timed behavioral sequence where it left off after the time-out period is over (Meck & Church, 1984; Meck, Church, Wenk, & Olton, 1987). On the other hand, a self-organizing avalanche circuit does have this competence (Grossberg, 1982, pp. 519–531; Grossberg & Kuperstein, 1986, Chap. 9). Moreover, each sensory representation in the avalanche can activate its own spectrally timed read-out to a drive representation.

This example illustrates the manner in which the totality of known temporally-discriminative neural networks have begun to delineate a global neural network architecture in which several distinct types of behavioral timing circuits cooperate to regulate the accurately timed autonomous unfolding of complex behaviors.

REFERENCES

- Berger, T. W., & Thompson, R. F. (1978). Neuronal plasticity in the limbic system during classical conditioning of the rabbit nictitating membrane response, I: The hippocampus. *Brain Research*, **145**, 323–346.
- Black, R. W., & Black, P. E. (1967). Heart rate conditioning as a function of interstimulus interval in rats. *Psychonomic Science*, **8**, 219–220.
- Boice, R., & Denny, M. R. (1965). The conditioned licking response in rats as a function of the CS-US interval. *Psychonomic Science*, **3**, 93–94.
- Burkhardt, P. E., & Ayres, J. J. B. (1978). CS and US duration effects in one-trial simultaneous fear conditioning as assessed by conditioned suppression of licking rats. *Animal Learning and Behavior*, **6**, 225–230.
- Cant, B. R., & Bickford, R. G. (1967). The effect of motivation on the contingent negative variation (CNV). *Electroencephalography and Clinical Neurophysiology*, **23**, 594.
- Carpenter, G. A., & Grossberg, S. (1987a). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, **37**, 54–115.
- Carpenter, G. A., & Grossberg, S. (1987b). ART 2: Stable self-organization of pattern recognition codes for analog input patterns. *Applied Optics*, **26**, 4919–4930.
- Carpenter, G. A., & Grossberg, S. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, **21**, 77–88.
- Cohen, M. A., & Grossberg, S. (1986). Neural dynamics of speech and language coding: Developmental programs, perceptual grouping, and competition for short term memory. *Human Neurobiology*, **5**, 1–22.
- Cohen, M. A., & Grossberg, S. (1987). Masking fields: A massively parallel architecture for learning, recognizing, and predicting multiple groupings of patterned data. *Applied Optics*, **26**, 1866–1891.
- Delacour, J., & Houcine, O. (1980). Conditioning to time: Evidence for a role of hippocampus from unit recording. *Neuroscience*, **23**, 87–94.
- Gormenzano, I., Kehoe, E. J., & Marshall, B. S. (1983). Twenty years of classical conditioning research with the rabbit. *Progress in Psychobiology and Physiological Psychology*, **10**, 197–275.
- Grossberg, S. (1971). On the dynamics of operant conditioning. *Journal of Theoretical Biology*, **33**, 225–255.
- Grossberg, S. (1972a). A neural theory of punishment and avoidance. I: Qualitative theory. *Mathematical Biosciences*, **15**, 39–67.
- Grossberg, S. (1972b). A neural theory of punishment and avoidance, II: Quantitative theory. *Mathematical Biosciences*, **15**, 253–285.
- Grossberg, S. (1975). A neural model of attention, reinforcement, and discrimination learning. *International Review of Neurobiology*, **18**, 263–327.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187–202.
- Grossberg, S. (1978). A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and

- plans. In R. Rosen and F. Snell (Eds.), *Progress in theoretical biology* (Vol. 5, pp. 233–374). New York: Academic Press.
- Grossberg, S. (1982). *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Boston: Reidel Press.
- Grossberg, S. (Ed.). (1987). *The adaptive brain, Volumes I and II*. Amsterdam: Elsevier/North-Holland.
- Grossberg, S. (Ed.). (1988). *Neural networks and natural intelligence*. Cambridge, MA: MIT Press.
- Grossberg, S., & Kuperstein, M. (1986). *Neural dynamics of adaptive sensory-motor control: Ballistic eye movements*. Amsterdam: Elsevier/North-Holland.
- Grossberg, S., & Levine, D. S. (1987). Neural dynamics of attentionally-modulated Pavlovian conditioning: Blocking, inter-stimulus interval, and secondary reinforcement. *Applied Optics*, **26**, 5015–5030.
- Grossberg, S., & Schmajuk, N. A. (1987). Neural dynamics of attentionally-modulated Pavlovian conditioning: Conditioned reinforcement, inhibition, and opponent processing. *Psychobiology*, **15**, 195–240.
- Grossberg, S., & Stone, G. O. (1986). Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review*, **93**, 46–74.
- Henneman, E. (1957). Relation between size of neurons and their susceptibility to discharge. *Science*, **26**, 1345–1347.
- Henneman, E. (1985). The size-principle: A deterministic output emerges from a set of probabilistic connections. *Journal of Experimental Biology*, **115**, 105–112.
- Hoechler, F. K., & Thompson, R. F. (1980). Effect of the inter-stimulus (CS-UCS) interval on hippocampal unit activity during classical conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*). *Journal of Comparative and Physiological Psychology*, **94**, 201–215.
- Irwin, D. A., Rebert, C. S., McAdam, D. W., & Knott, J. R. (1966). Slow potential change (CNV) in the human EEG as a function of motivational variables. *Electroencephalography and Clinical Neurophysiology*, **21**, 412–413.
- Jones, R., & Keck, M. J. (1978). Visual evoked response as a function of grating spatial frequency. *Investigative Ophthalmology and Visual Science*, **17**, 652–659.
- Killeen, P. R., & Weiss, N. A. (1987). Optimal timing and the Weber function. *Psychological Review*, **94**, 455–468.
- Konorski, J. (1948). *Conditioned reflexes and neuron organization*. London: Cambridge University Press.
- Low, M. D., Borda, R. P., Frost, J. D., & Kellaway, P. (1966). Surface negative slow potential shift associated with conditioning in man. *Neurology*, **16**, 711–782.
- Luce, E. S., & Fuchs, A. F. (1972). Activity of brain stem neurons during eye movements of alert monkeys. *Journal of Neurophysiology*, **35**, 445–461.
- Maricq, A. V., Roberts, S., & Church, R. M. (1981). Methamphetamine and time estimation. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 18–30.
- McAdam, D. W. (1969). Increases in CNS excitability during negative cortical slow potentials in man. *Electroencephalography and Clinical Neurophysiology*, **26**, 216–219.
- McAdam, D., Knott, J. R., & Chiorini, J. (1965). Classical conditioning in the cat as a function of the CS-US interval. *Psychonomic Science*, **3**, 89–90.
- Meck, W. H. (1983). Selective adjustment of the speed of internal clock and memory processes. *Journal of Experimental Psychology: Animal Behavior Processes*, **9**, 171–201.
- Meck, W. H., & Church, R. M. (1984). Simultaneous temporal processing. *Journal of Experimental Psychology (Animal Behavior)*, **10**, 1–29.
- Meck, W. H., & Church, R. M. (1987). Cholinergic modulation of the content of temporal memory. *Behavioral Neuroscience*, **101**, 457–464.
- Meck, W. H., Church, R. M., Wenk, G. L., & Olton, D. S. (1987). Nucleus basalis and magnocellularis and medial septal area lesions differentially impair temporal memory. *Journal of Neuroscience*, **7**, 3505–3511.
- Millenson, J. R., Kehoe, E. J., & Gormenzano, I. (1977). Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS-US intervals. *Learning and Motivation*, **8**, 351–366.
- Musselwhite, M. J., & Jeffreys, D. A. (1985). The influence of spatial frequency on the reaction times and evoked potentials recorded to grating pattern stimuli. *Vision Research*, **25**, 1545–1555.
- Parker, D. M., & Salzen, E. A. (1977a). Latency changes in the human visual evoked response to sinusoidal gratings. *Vision Research*, **17**, 1201–1204.
- Parker, D. M., & Salzen, E. A. (1977b). The spatial selectivity of early and late waves within the human visual evoked response. *Perception*, **6**, 85–95.
- Parker, D. M., Salzen, E. A., & Lishman, J. R. (1982a). Visual-evoked responses elicited by the onset and offset of sinusoidal gratings: Latency, waveform, and topographic characteristics. *Investigative Ophthalmology and Visual Science*, **22**, 675–680.
- Parker, D. M., Salzen, E. A., & Lishman, J. R. (1982b). The early waves of the visual evoked potential to sinusoidal gratings: Responses to quadrant stimulation as a function of spatial frequency. *Electroencephalography and Clinical Neurophysiology*, **53**, 427–435.
- Plant, G. T., Zimmern, R. L., & Durden, K. (1983). Transient visually evoked potentials to the pattern reversal and onset of sinusoidal gratings. *Electroencephalography and Clinical Neurophysiology*, **56**, 147–158.
- Rawlins, J. N. P. (1985). Associations across time: The hippocampus as a temporary memory store. *The Behavioral and Brain Sciences*, **8**, 479–496.
- Rawlins, J. N. P., Feldon, J., & Gray, J. A. (1982). Behavioral effects of hippocampectomy depend on inter-event intervals. *Society for Neuroscience Abstracts*, **8**, 22.
- Robinson, D. A. (1970). Oculomotor unit behavior in the monkey. *Journal of Neurophysiology*, **35**, 393–404.
- Schiller, P. H. (1970). The discharge characteristics of single units in the oculomotor and abducens nuclei of the unanesthetized monkey. *Experimental Brain Research*, **10**, 347–362.
- Schneiderman, N. (1972). Response system divergencies in aversive classical conditioning. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning, II: Current research and theory*. New York: Appleton Century Crofts.
- Skrandies, W. (1984). Scalp potential fields evoked by grating stimuli: Effects of spatial frequency and orientation. *Electroencephalography and Clinical Neurophysiology*, **58**, 325–332.
- Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, **3**, 679–687.
- Solomon, P. R. (1979). Temporal versus spatial information processing views of hippocampal functions. *Psychological Bulletin*, **86**, 1272–1279.
- Solomon, P. R. (1980). A time and a place for everything? Temporal processing views of hippocampal function with special reference to attention. *Physiological Psychology*, **8**, 254–261.
- Solomon, P. R., van der Schaaf, E. R., Thompson, R. F., & Weisz, D. J. (1986). Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behavioral Neuroscience*, **100**, 729–744.
- Staddon, J. E. R. (1983). *Adaptive behavior and learning*. Cambridge, England: Cambridge University Press.
- Vassilev, A., Manahilov, V., & Mitov, D. (1983). Spatial frequency and pattern onset-offset response. *Vision Research*, **23**, 1417–1422.

- Vassilev, A., & Strashimirov, D. (1979). On the latency of human visually evoked response to sinusoidal gratings. *Vision Research*, *19*, 843–846.
- Walter, W. G. (1964). Slow potential waves in the human brain associated with expectancy, attention, and decision. *Arch. Psychiat. Nervenkr.*, *206*, 309–322.
- Wilkie, D. M. (1987). Stimulus intensity affects pigeons' timing behavior: Implications for an internal clock model. *Animal Learning and Behavior*, *15*, 35–39.
- Williamson, S. J., Kaufman, I., & Brenner, D. (1978). Latency of the neuromagnetic response of the human visual cortex. *Vision Research*, *18*, 107–110.