COMPUTATIONAL AND BEHAVIORAL INVESTIGATIONS OF REAL-TIME MODELS OF CLASSICAL CONDITIONING

Diana E. J. Blazis

COINS Technical Report 90-46
June 4, 1990

University of Massachusetts
Amherst, Massachusetts 01003

This research was supported by Air Force Office of Scientific Research grant AFOSR 86-0182, by National Science Foundation grant NSF BNS 88-10624, and by a University of Massachusetts Graduate Fellowship.

COMPUTATIONAL AND BEHAVIORAL INVESTIGATIONS OF REAL-TIME MODELS OF CLASSICAL CONDITIONING

A Dissertation Presented

by

DIANA E.J. BLAZIS

Submitted to the Graduate School

of the University of Massachusetts in partial fulfillment

of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 1990

Neuroscience and Behavior

COMPUTATIONAL AND BEHAVIORAL INVESTIGATIONS OF REAL-TIME MODELS OF CLASSICAL CONDITIONING

A Dissertation Presented

by

DIANA E.J. BLAZIS

Approved as to style and content by:

John W. Moore, Chairman

Andrew G. Barto, Member

Jeffrey D. Blaustein, Member

Gordon A. Wyse, Member

Rodney K. Murphey, Graduate Program Director

Graduate Program in Neuroscience and Behavior

DEDICATIONS

In loving memory of my Dad,
John H. Jennings,
and to
Scott, Mom, Nancy, and Gary,
without whom I would not have
travelled this far.

ACKNOWLEDGEMENTS

It is a pleasure and an honor to thank the many people who have contributed to my education and to this dissertation.

First, I thank John W. Moore, who has provided me with expert guidance, research opportunities, and support for many years. I thank Andrew Barto and Richard Sutton, whose ideas provided a framework for my research in learning, for their frequent assistance with theoretical issues.

To my thesis committee, Andrew G. Barto, Jeffrey D. Blaustein, Gordon A. Wyse, and John Moore, I express my gratitude for useful comments and encouragement. To John go special thanks for his editorial assistance.

I am grateful for advice and technical and moral support from my colleagues Neil Berthier, John "JD" E. Desmond, William G. Richards, Thomas N. Ricciardi, and Marcy Rosenfield. All of these individuals contributed to the construction of the behavioral apparatus, which was designed by Dr. Richards. Drs. Richards and Desmond shared their computing expertise during the development of my data analysis software. Dr. Berthier maintained the computer network. Tom Ricciardi expedited the end of this project by assisting in the conduct of Experiment 4. Marcy Rosenfield provided dedicated care of all the animals used in this research.

My family and friends were a constant source of love and encouragement, and for this I deeply thank them all. Most importantly, I thank my husband Scott, for his love and support, his faith in my abilities, and for sacrifices large and small.

ABSTRACT

COMPUTATIONAL AND BEHAVIORAL INVESTIGATIONS
OF REAL-TIME MODELS
OF CLASSICAL CONDITIONING

MAY 1990

DIANA E.J. BLAZIS

B.S., UNIVERSITY OF MASSACHUSETTS

PH.D., UNIVERSITY OF MASSACHUSETTS

Directed by: Professor John W. Moore

Real-time models of classical conditioning were evaluated using simulation and behavioral techniques. Performance criteria included ability to account for aspects of the conditioned response (CR) during conditioning of the rabbit nictitating membrane response (NMR), such as development of CR topography and the dependency of conditioning on the interstimulus interval (ISI) elapsing between onsets of conditioned and unconditioned stimuli (CSs and USs).

The models investigated included the Sutton-Barto-Desmond (SBD) model (Moore, Desmond, Berthier, Blazis, Sutton, and Barto, 1986), which predicts CR topography, and the Temporal Differences (TD) model (Sutton and Barto, 1987), which does not. Simulation studies addressed the impact of modifying the TD model to describe topography. The resultant model, called the TD_{RTS} model, replicates the learning phenomena described by the TD model but is highly parameter-sensitive. Simulation studies showed the performance of the TD, TD_{RTS}, and SBD models to

be similar in most respects. However, model predictions differ for simultaneous and backwards conditioning protocols involving single and multiple CSs.

Behavioral verification of new predictions of the models was carried out using rabbit NMR conditioning. Experiment 1 showed that with fixed-trace-interval conditioning, conditioning increases with CS duration, a result consistent with the SBD but not the TD model. Experiment 1 also showed a novel CS intensity by ISI interaction that none of the models explain. Experiment 2 replicated Experiment 1 and analysed response topography to illustrate problems caused by the CS representation assumed by the models. Experiments 3 and 4 investigated compound conditioning using one CS followed by simultaneous presentation of a second CS and US. Experiments 3 and 4 showed that conditioning to the compound occurred more readily than in controls when the first CS was less salient than the second. Experiment 3 presented weak evidence that the simultaneous CS was inhibitory. Of the models explored, the TD model most fully accounts for the results of Experiments 3 and 4.

These computational and behavioral experiments suggested that the models might be better served by alternative CS representations, and highlighted the need for a resolution of issues in simultaneous conditioning. These points and future directions are discussed.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Real-time Models of Conditioning

Approaches to mathematical modeling of conditioning have evolved from a relatively coarse description of the trial-level aspects of conditioning (e.g., Hull, 1943; Rescorla and Wagner, 1972; Pearce and Hall, 1980) to a more detailed description of the intra-trial aspects of learning and of the conditioned response or CR; models that provide such detailed descriptions are sometimes referred to as *real-time models*. Figure 1.1 illustrates the distinction between real-time and trial-level models of conditioning. The term V_i is used as an index of the conditioned strength of the i^{th} conditioned stimulus or CS. Learning is defined as changes in V_i, and designated as ΔV_i.[1] Trial-level models compute changes in V_i only once per trial, but real-time models allow for the possibility that changes in V_i are computed during particular

[1]The usage of the term V_i to denote associative strength is common in psychology, continuing the tradition established by the Rescorla-Wagner model (Rescorla and Wagner, 1972). In the context of the present research, V_i is assumed to map directly onto percent CRs, a behavioral measurement typically used to determine strength of conditioning. Notably, this assumption leads to the prediction of negatively accelerated learning curves (Moore et al., 1986; Sutton and Barto, 1981; Sutton and Barto, 1987). In most conditioning situations, learning curves tend to be sigmoidal. For further discussion of this point, see Chapter 2.

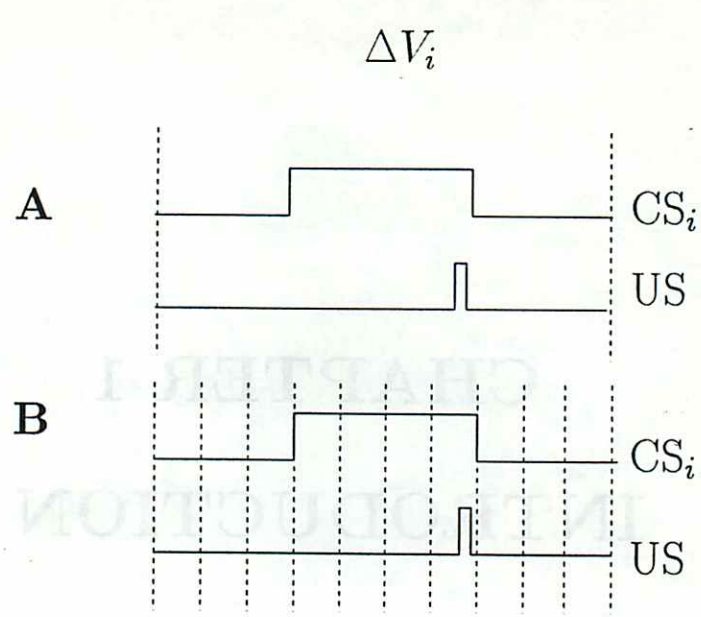

Figure 1.1: **Real-time and Trial-level Models of Conditioning.** V_i is an index of the conditioned strength of the i^{th} CS. Learning is defined as changes in V_i, and designated as ΔV_i. The solid lines denote onsets and offsets of stimuli. Dotted verticals depict the boundaries of computational epochs or time steps. Trial-level models compute changes in V_i only once per trial, as represented in A. Real-time models compute changes in V_i during time steps within a trial, as represented in B.

time steps or time bins that are smaller than the trial duration. In Figure 1.1, changes in V_i can occur before, during, and after stimuli are presented.[2]

The evolution of real-time models has been driven by the need to address certain temporal aspects of conditioning, and has been facilitated by the arrival of powerful and readily available computing technology. Real-time models include those developed by Desmond and Moore (1988), Grossberg and Schmajuk (1989), Klopf (1988), Moore, Desmond, Sutton, Barto, Blazis, and Berthier (1986), Moore and Stickney (1980, 1985), Schmajuk and Moore (1988), and Sutton and Barto (1981).

Because changes in V_i might occur throughout a classical conditioning trial, real-time models can potentially account for important characteristics of classical conditioning:

[2]It should be noted that to the mathematician *real-time* denotes continuous computations. The models discussed in the present research are correctly known as *discretized real-time* models, but it has been customary to refer to them as simply real-time models (e.g., Klopf, 1988).

• **The strength of classical conditioning depends on the temporal arrangement of stimuli.** Strength of conditioning is dependent upon the duration of the interval between stimulus onsets, the *interstimulus interval* or ISI. The *ISI function* is typically plotted as percentage CRs obtained after an arbitrary number of trials against ISI duration. The typical ISI function is concave downward, rising sharply from a minimal ISI to peak at a value denoted the *optimal ISI*, then declining as the ISI increases beyond the optimal range. The ISI function is only one example of the temporal sensitivity of conditioning. Other examples can be found in the literature involving conditioning of multiple CSs (Kehoe, 1982).

• **The CR is an anticipatory response.** The time course or *topography* of the CR has notable features. Over training, onset latency of the CR decreases and response amplitude increases. After a sufficient number of CS and unconditioned stimulus (US) pairings, the initiation of the CR occurs before the onset of the US (see Gormezano, Kehoe, and Marshall, 1983). In many training procedures peak amplitude of the response occurs at or near the temporal locus of the US. These general features of the CR can be observed not only at the behavioral level, but may be observed at the level of single cells (see for example, Berthier and Moore, 1986; Desmond and Moore 1986; Berger, Rinaldi, Weisz, and Thompson, 1983).

The temporal domain of real-time models allows them to address issues of stimulus timing and response topography, and also suggests ways in which models of learning could be implemented at the level of neurons (Barto and Sutton, 1982; Blazis and Moore, 1987). Hence, mathematical models of classical conditioning can link up with neurophysiological investigations of the CR.

1.2 The Conditioned Rabbit Nictitating Membrane Response

One conditioning preparation used to evaluate real-time models is the conditioned rabbit nictitating membrane response or NMR. The rabbit NMR is a protective response resulting from retraction of the eyeball and the passive sweeping of the NM over the eye (Berthier, 1984; Berthier and Moore, 1980). Reports from studies of the NMR preparation date back to the 1960's (Gormezano, Schneiderman, Deaux, and Fuentes, 1962). The literature shows that the temporal dependencies of classical conditioning (e.g., ISI effects and anticipatory nature of the response) are manifested by the NMR, that many conditioning phenomena can be shown by this response (Gormezano, Kehoe, and Marshall, 1983), and that response topography can be complex (Millenson, Kehoe, and Gormezano, 1977). In addition, the NMR preparation itself is widely used for anatomical and physiological studies of learning.

Real-time models developed with the rabbit NMR system in mind can be used to design new experiments. This dissertation uses the NMR as both database for and testbed of models of learning.

1.3 Review of Selected Real-Time Models

1.3.1 Hebb's Postulate

Many real-time models can trace their ancestry to Hebb's postulate, an early attempt to relate learned changes in behavior to cellular mechanisms. Hebb (1949) suggested that learning can be conceptualized as a change in synaptic (or connection) strength that results from the conjunction of post- and pre-synaptic activity. When presynaptic activity is followed directly by postsynaptic activity, the connection

between pre- and post-synaptic elements is strengthened[3]Postsynaptic activity at time step t is denoted as $s(t)$:

$$s(t) = f\left(\sum_{i=1}^{n} V_i(t)x_i(t)\right), \tag{1.1}$$

where the magnitude of presynaptic activity elicited by the i^{th} stimulus at time t is denoted $x_i(t)$. Each stimulus input i, $i = 1, \ldots, n$, has a variable synaptic weight or efficacy whose value at a given time t is denoted $V_i(t)$. Learning, denoted as δV_i, is equal to $V_i(t) - V_i(t-1)$, and is computed as:

$$\Delta V_i(t) = cx_i(t)s(t). \tag{1.2}$$

The constant c specifies the learning rate.

Equations 1.2 and 1.1 can be applied to the classical conditioning situation by assuming that single neuron-like units show simplified analogs of animal behavior in conditioning experiments (Barto and Sutton, 1982; Klopf, 1972; Sutton and Barto, 1987). In this context, presynaptic activity is aligned with the input of the i^{th} stimulus, x_i in Equations 1.2 and 1.1. "Stimulus" refers to either the CS_i or the US, but it is often assumed that although US inputs contribute to s, the synaptic weight for the US is fixed (e.g., Sutton and Barto, 1981). Postsynaptic activity, s, is a component of or a contributing signal to the CR or unconditioned response (UR), and is computed by Equation 1.1. Changes in V_i are assumed to occur at units that receive convergent inputs of both the US and CS_i. At the start of conditioning, V_i is assumed to be 0. The weight of the US is generally assumed to be positive. Thus, initial pairings of the US and CS_i result in an output that is due solely to activity on the US input pathway. Continued pairings of the CS_i and US increase V_i to the point at which s also includes a contribution from activity elicited by the CS_i: the CR.

[3]Equation 1.2 is one kind of Hebbian *rule*, any model that takes the general form $\Delta V = f(g(x)h(y))$. Hebb did not supply a mathematical equation describing his postulate.

Though the model shown in Equation 1.2 may be consistent with some kinds of learning, without additional assumptions it cannot account for basic facts of conditioning. For example, by Equation 1.2 optimal conditioning would be obtained when the CS and US begin at the same time (ISI = 0). However, the typical ISI function peaks at an ISI greater than zero. As discussed by Sutton and Barto (1981), the specification of suitable delays between the CS and US pathways can extend Hebb's model to the description of appropriate ISI functions, but the cost is that CR onset latency is delayed to the point of US onset.

1.3.2 The Sutton-Barto (SB) Model

Sutton and Barto (1981) modified Hebb's model by assuming that learning is affected by the coincidence of a non-stimulating trace of presynaptic activity with *changes* in post-synaptic activity. The Sutton and Barto (SB) model assumes that conditioning occurs to the extent that the *prediction* of levels of stimulation differs from levels that actually occur. Therefore, classical conditioning is the process of learning to anticipate changes in levels of stimulation.

In the SB model, changes in the associative strength of CS_i, at time t, denoted as $\Delta V_i(t)$, are computed as follows:

$$\Delta V_i(t) = c[s(t) - \bar{s}(t)]\bar{x}_i(t), \tag{1.3}$$

where c is a learning rate parameter, $0 < c \leq 1$, $s(t)$ is the output at time step t, $\bar{x}_i(t)$, defined below, is the eligibility of CS_i, and $\bar{s}(t)$, the prediction of future output levels, is defined as:

$$\bar{s}(t) = \beta\bar{s}(t-1) + (1-\beta)s(t-1) \tag{1.4}$$

where $0 < \beta < 1$. In Sutton and Barto (1981), $\bar{s}(t) = s(t-1)$, the special case which results when $\beta = 0$. The eligibility trace, a non-stimulating *stimulus trace*, is defined as follows:

$$\bar{x}_i(t) = \alpha \bar{x}_i(t-1) + x_i(t-1) \tag{1.5}$$

where $0 \leq \alpha < 1$, and x_i, the input of CS_i, is a rectangular pulse equal to 1 when the CS is on and 0 when the CS is off.[4] trace that does not influence the level of s (Equation 1.1). The function of $\bar{x}$ becomes clear by inspection of Equation 1.3. The coincidence of pre-synaptic and post-synaptic activity does not automatically affect V_i, but instead determines the extent to which V_i may be modified. Since the connection does not become eligible for modification until some time after CS onset, increments in V_i do not occur unless the US occurs after the onset of the CS. Thus, $\bar{x}$ allows the SB model to account for the fact that the most robust conditioning is typically obtained with an ISI that is greater than zero.

Since conditioning can occur when the onset of the CS precedes that of the US, given a sufficient number of trials the output, s, can be greater than zero prior to the onset of the US. Thus, the SB model provides an account of the anticipatory nature of the CR.

In summary, the SB model specifies that changes in associative strength V_i occur when (a) the connection is eligible for modification and (b) there are changes in the output over time steps. The SB model accounts for a wide variety of conditioning phenomena, including acquisition, extinction, and ISI functions and the anticipatory nature of the CR. The model makes interesting predictions for a conditioning

[4]Sutton and Barto (1981) provided a lengthy discussion about the eligibility or stimulus trace, a notion borrowed from Klopf (1972). The stimulus trace can be either stimulating or non-stimulating, where the former affects s but the latter does not. Sutton and Barto (1981) noted the fact that the eligibility trace implements the idea that a stimulus activates an internal trace that can persist beyond stimulus offset, as put forth by Pavlov (1927), Hull (1943), and others. The internal trace could arise via reverberatory activity, that is, some type of stimulating trace, or it could arise via non-electrical mechanisms such as the concentration of chemicals or ions. Sutton and Barto advocated the idea that a non-stimulating

procedure called blocking[5]. This prediction has been tested and generally supported using the conditioned rabbit NMR (Kehoe, Schreurs, and Graham, 1987). However, the SB model does fail to predict behaviorally appropriate results in some instances (Klopf, 1988; Moore et al., 1986; Sutton and Barto, 1987). For example, the SB model does not yield appropriate ISI functions when realistically short US durations are used.

Although the SB model provides a useful algorithm for generating anticipatory CRs, the topography of the CR that results is not correct. Without additional machinery, CRs generated by the SB model emerge with zero latency and with peak amplitude at CS onset. In contrast, a real conditioned response like that obtained with the rabbit NMR begins well after CS onset and rises gradually in a ramped or S-shaped fashion within the ISI. The CR attains a maximum at or near the temporal locus of the US, and then decays rapidly during the post-US period. This pattern of response topography is also reflected in the activity of some types of neurons that have been identified in single-unit recording studies as being linked to the generation of CRs (Desmond, 1985; see also Desmond and Moore, 1986).

1.3.3 The Sutton-Barto-Desmond (SBD) Model

Desmond, Moore, and colleagues parameterized the SB model, constraining it to produce reasonable approximations of ISI functions and of response topography for forward-delay conditioning (Moore et al., 1986). Parameter values were set at those that produced the best joint description of these aspects of NMR conditioning. Then, the model's performance was evaluated for a variety of classical conditioning procedures. The resulting model, the Sutton-Barto-Desmond (SBD) model, retained and improved upon the SB model's ability to describe classical conditioning, while

[5]In the Kamin blocking paradigm (1968), one CS is conditioned and then paired with another CS. The compound of the novel and the pretrained CSs is reinforced, that is, paired with a US. Generally, very little associative strength accrues to the added CS.

extending the SB model to the simulation of NMR topography and conditioning-related neural firing.

The SBD model retains the learning rule of the SB model, but incorporates other computational mechanisms to predict appropriate results for classical conditioning of the rabbit NMR:

- In order to predict conditioning with short USs, the SBD model assumes a variable US effectiveness term, λ'. Essentially, the effect of the US decreases as conditioning (V_i) increases.
- To account for ISI functions, the SBD model assumes that the decay of the eligibility trace depends on the CS duration. The eligibility of a CS for changes in V_i declines after CS offset, but the rate of decay of eligibility increases with the CS duration. This feature of the SBD model is discussed more fully in Chapter 4.
- In order to generate features of response topography, the SBD model assumes that the input or *representation* of CS_i, x_i, is a *template* for the desired response. The input x_i rises in an S-shaped fashion while CS_i is on. At CS_i offset, x_i decays geometrically to 0. Because of the assumptions made about the CS representation, the SBD model predicts that CR latency decreases and CR amplitude increases as conditioning proceeds.

The SBD model is presented in its entirety in Appendix A.

In modeling response topography, the developers of the SBD model sought to combine a reasonably successful model of classical conditioning with physiological investigations of the rabbit conditioned NMR.[6] Analysis of the restricted set of

[6]Some contemporary learning theorists (see Rescorla, 1988) disagree with the tactic of closely observing the topography of the CR, arguing that it may not appropriate to infer that learning has

parameter values required for optimum performance of the SBD model led to the specification of neural circuits which might be capable of computing changes in associative strength as specified by the SB model (Blazis and Moore, 1987). Specifically, Blazis and Moore described how changes in associative strength, as defined in the SB model, could be computed by circuits in the cerebellar cortex, a region thought to be essential for the learning and performance of the conditioned rabbit NMR (Thompson, Donegan, Clark, Lavond, Lincoln, Madden, Mamounas, Mauk, and McCormick, 1987; Yeo, Hardiman, Glickstein, 1985a-c, 1986).

1.3.4 The Temporal Differences (TD) Model

The SBD model provides one approach to the solution of some of the problems of the SB model. However, Sutton and Barto (1987) argue that the model does so at the expense of computational simplicity as two new variables, the US effectiveness term and variable eligibility decay indicated above, are added. In order to encompass conditioning with a short US and describe ISI effects, Sutton and Barto (1987) propose instead the Temporal Differences (TD) model.

Like the SB and SBD models, the TD model views conditioning as a process of prediction. The TD model specifies that in conditioning the animal (or neuron-like unit) predicts a discounted sum of future occurrences of the US signal (Sutton and Barto, 1987). Put another way, the quantity to be predicted is the area of a discounted US. Discounting of future predictions of the US is assumed in order to

occurred simply because there is a change in a CS's ability to evoke a response, and that reliance upon a particular CR can blur the distinction between learning and performance. For example, different CSs, when trained with the same US, can evoke profoundly different responses. Rescorla (1988) suggests other ways to detect learning, such as observing the ability of the CS to disrupt conditioning to other stimuli (blocking) or to reinforce learning about other CSs (higher-order conditioning). Such testing procedures may not be compatible with the requirements of many neurophysiological model systems. Hence, it may be more appropriate for the neurophysiologist to regard Rescorla's criticisms as an impetus to include proper control procedures when necessary and to select carefully the response system to be observed.

account for the fact that conditioning tends to decrease as ISI increases. Increases and decreases in V_i are in part due to the time difference between imminent and future values of predictions about the US signal.

Unlike the SB and SBD models, the TD model assumes that the US is involved directly in weight changes. In the TD model, the US is modeled by a signal whose value at time t is denoted as $\lambda(t)$. It takes a value > 0 if the US is present at time t and the value 0 if the US is absent at time t. In simulation experiments described in Sutton and Barto (1987), each time step corresponds to 50 ms of real time.

$\Delta V_i(t)$, is computed as follows:

$$\Delta V_i(t) = c(\lambda(t) + \gamma P(V(t), x(t)) - P(V(t), x(t-1)))\bar{x}_i(t), \tag{1.6}$$

where $c > 0$ and $0 < \gamma < 1$. The parameter γ in Equation 1.6 discounts future predictions for the value of the US signal λ. In doing so, this parameter allows a CS_i with existing associative strength to generate changes in V_i throughout its presence. Variable $\lambda(t)$ is the value of a signal indicating the presence of the US and $P(V, x)$ is defined as

$$P(V, x) = \begin{cases} \sum_{i=1}^{n} V_i x_i, & \text{if } \sum_{i=1}^{n} V_i, x_i > 0; \\ 0, \text{ otherwise.} \end{cases} \tag{1.7}$$

where V_i is associative strength and x_i is the value of the input of the i^{th} CS. P can be interpreted as a prediction about the future value of the US signal, λ. Sutton and Barto (1987) assumed that CS_i input is equal to 1 when CS_i is present and 0 when CS_i is not present.

Variable $\bar{x}_i(t)$ in Equation 1.6 is the input pathway eligibility of CS_i at time step t:

$$\bar{x}(t) = \beta\bar{x}(t-1) + (1-\beta)x(t-1), \tag{1.8}$$

The effects of stimulus presences, onsets, and offsets are specified somewhat differently in the TD and SBD learning models. In both, changes in V_i occur during

CS onsets and offsets. However, in the TD model, the presence rather than onset and offset of the US causes changes in synaptic weights (Equations 1.6 and 1.7). This feature alleviates one problem of the SB model, that being that USs must be unrealistically long for conditioning to occur (Sutton and Barto, 1987).

The TD model encompasses a greater range of conditioning data than does the SB model. Appropriate ISI functions are obtained for both forward-delay and trace conditioning procedures. The model is capable of describing many other conditioning phenomena (Sutton and Barto, 1987). However, the model has not been parameterized to generate detailed response topography. Chapter 2 details issues of modeling response topography with the TD model.

1.4 Initial Comparison of the TD and SBD Models

The TD and SBD models represent two different approaches to modeling classical conditioning. The TD model is based in an engineering theory of prediction (this point is discussed further in Chapter 2) and is geared towards describing general behavioral results. The SBD model represents a hybrid between prediction algorithm and mathematical devices that, while biologically based, complicate mathematical analysis, and it is specifically geared towards predicting the results of conditioning of the rabbit NMR. However, both models share common predictions of a wide range of behaviors. This section indicates their common traits, their differential predictions, and their shortcomings.

1.4.1 Differential Predictions of the TD and SBD Models

Both the SBD and TD models describe a great range of NMR conditioning data. Phenomena that are readily described by both models include many aspects of

single-CS conditioning, such as acquisition and extinction for both forward-delay and trace conditioning. Multiple-CS phenomena like higher-order conditioning, blocking, and conditioned inhibition are predicted by both models. However, notable differences exist between the models that allow the specification of experiments to test basic features of each model. In this section, critical areas of disagreement between the models' predictions are briefly presented.

Simultaneous and Backwards Conditioning

With the rabbit NMR preparation, ISIs greater than zero (CS onset preceding US onset) typically result in excitatory conditioning. However, there remains no agreement about the type of conditioning obtained at a simultaneous ISI (simultaneous onset of CS and US) or a backwards ISI (US onset precedes CS onset).[7] CSs trained at these ISIs can potentially have associative strengths that are neutral (that is, 0 in the modeling interpretation), inhibitory (less than 0 in the modeling interpretation), or excitatory (greater than 0).

The experimental evidence suggests that simultaneous and backwards conditioning can be excitatory or inhibitory, depending on the response system and tasks used for assessment. In the conditioned suppression literature, there is evidence for both inhibitory (for example, Heth, 1976) and excitatory conditioning (for example, Burkhardt and Ayres, 1978). To date, few studies with the NMR preparation directly address whether a simultaneously conditioned CS is inhibitory or excitatory. The studies that do exist (for example, Kehoe, Feyer and Moses, 1981; Schneiderman, 1966; Smith, Coleman and Gormezano, 1969) suggest that a simultaneously conditioned CS is neutral or inhibitory. Backward conditioning of the rabbit NMR

[7] In a simultaneous or backwards conditioning procedure, the elicitation of the UR during paired CS and US trials obscures possible CRs, so CS-alone or *probe* trials are presented. In order to assess inhibition or excitation, *retardation* tests, in which the animal receives CS pairings at a forward ISI, can be conducted, with rate of learning being assessed against control subjects that have not received simultaneous or backwards conditioning.

is thought to be inhibitory, as evidenced by retarded acquisition of a forward ISI following training with a backwards ISI (Plotkin and Oakley, 1976).

Recent studies may provide a reconciliation of the long and controversial history of simultaneous associations. Based on data from a series of conditioned-lick suppression studies in rats, Matzel, Held, and Miller (1988) maintained that simultaneous associations are learned, but that animals responded in a way that reflected that information only "when the information was of functional, predictive value appropriate for the response system ... monitored" (p. 341). Tait and Saladin (1986) showed that the same backwards conditioning task in the rabbit NMR could facilitate subsequent learning of subsequent conditioned punishment task yet retard learning of a forward-delay task with the rabbit NMR. Based on these data the investigators suggested that CS-US associations have a multi-dimensional structure. Computational studies of adaptive networks have also supported the idea that associations have a dual character: a linkage to the US, and an expectation of when the US will occur (Desmond and Moore, 1988).

The TD and SBD models reflect the controversy in the field by predicting differential results for simultaneous and backwards conditioning. The TD model predicts excitatory simultaneous conditioning (Sutton and Barto, 1987), while the SBD model predicts inhibitory conditioning (Moore et al., 1986). For backwards conditioning, the TD model predicts that no learning will occur, but the SBD model again predicts inhibitory conditioning. These predictions have implications, not only for single-CS conditioning, but also in compound conditioning procedures where one CS is followed by a second CS that occurs at the same time as the US. Results of simulation and behavioral investigation of the discrepant model predictions are presented in Chapter 4.

Trace Conditioning

The SBD and TD models rely on different mechanisms to predict the form of the concave-downward ISI function. These mechanisms are described and compared in Chapter 4. The present discussion is concerned with the form of the ISI function in trace conditioning. In trace conditioning, US onset occurs some time after CS offset. The period of time elapsing between CS offset and US onset is called the *trace interval.* The TD model predicts that regardless of the size of the trace interval, conditioning occurs more readily with shorter ISIs than with longer ISIs. However, the SBD model predicts that, given a sufficiently long and fixed trace interval, conditioning occurs more readily with longer-duration CSs (that is, longer ISIs) than short-duration CSs (that is, shorter ISIs).. This differential prediction of the models is investigated through simulation and behavioral verification in Chapter 5.

1.4.2 Shortcomings of TD and SBD Models

Both the TD and SBD models share shortcomings in their ability to describe certain conditioning phenomena. Neither model can predict spontaneous recovery or rapid reacquisition of an extinguished conditioned response. To encompass these phenomena, additional assumptions would have to be made, for example, that V_i, once trained, does not completely extinguish to a value of 0. Sutton and Barto (1987) pointed out that certain phenomena in conditioning with multiple CSs are also not addressed by the models. For example, the TD and SBD models do not predict results of sensory preconditioning experiments[8] (Rescorla, 1980). Sutton and Barto (1987) suggested that sensory preconditioning may be mediated by so-called "indirect associations", that is, associations among the stimuli as well as between

[8] A sensory preconditioning has at least three stages. In the first stage, CS1 is paired with CS2. In the second stage, CS2 is paired with a US. In the third stage, CS1 is tested. Responding to CS1 in the third stage is interpreted as reflecting an association between CS1 and CS2 (Rescorla, 1980).

these stimuli and the US. The TD and SBD models make no predictions about such associations, and are therefore incapable of predicting phenomena such as sensory preconditioning. However, the models could in principle be modified to encompass these results (Moore et al., 1986; Sutton and Barto, 1987). Computational and behavioral studies of these models, such as the ones in this dissertation, assess the utility of executing such modifications.

1.5 Research Objectives and Reader's Guide

Figure 1.2 illustrates a general approach used by Moore and colleagues to align models of classical conditioning with neural substrates involved in the classically conditioned NMR of the rabbit. The approach consists of applying increasingly more stringent data constraints to a model and assessing the consequences in simulation experiments. A model so constrained can be aligned with aspects of NMR anatomy and physiology. Throughout the process there exists a potential for new experiments at behavioral and physiological levels.

The research presented in this dissertation is a continuation of the above research strategy for the SBD model. It also applies the strategy to the TD model in an effort to compare the two as descriptions of classical conditioning. Since the TD model as described by Sutton and Barto (1987) does not encompass response topography, one objective of the present research was to modify or extend the TD model to allow it to capture basic features of the CR. The present research includes behavioral experiments which assess predictions of the SBD model. Since the TD and SBD models predict different outcomes for the experiments reported here, the behavioral experimentation had consequences for the TD model as well.

The next two chapters describe computational studies of the TD model. Chapter 2 establishes the methodology and set of assumptions for interpreting model

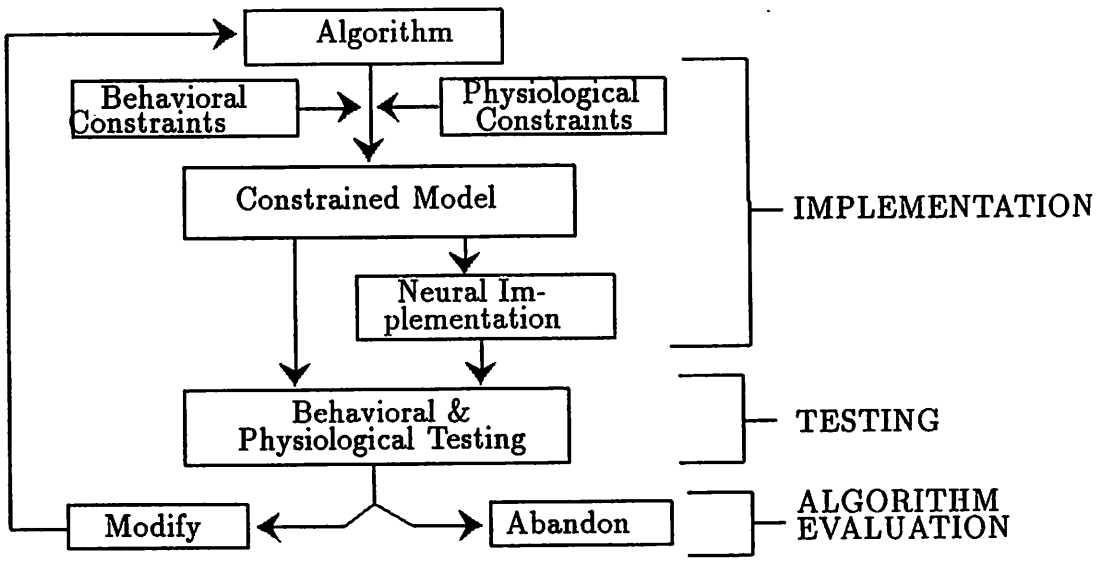

Figure 1.2: A Research Strategy for Developing Computational Models of Classical Conditioning. A general real-time model of classical conditioning is constrained by existing physiological data; key features of a model, for example, V_i, outputs, and inputs, are specified so that the model can generate appropriate response topographies. Performance of the model is compared with the behavioral literature. Such an assessment may or may not generate additional constraints. Finally, the constrained model is aligned with aspects of NMR anatomy and physiology. Models designed via this process may generate novel behavioral and neurophysiological predictions which can be subjected to empirical verification in the Testing stage. Experiments can direct the modification of certain parameters and variables, or disprove the model as a description of classical conditioning.

variables like associative strength V_i. Additionally, Chapter 2 considers issues of modeling realistic response topography with the TD model, and suggests modifications of both the TD learning rule and the form of CS input. The goal of the modeling work was that predicted response topography be consistent with what is observed during classical conditioning of the rabbit NMR. By this criterion, and to facilitate comparisons with the SBD model, CS representation in simulation experiments with the TD model derived from the provision of an input trace that is a template for the desired response. Chapter 3 describes simulation experiments investigating parameter values that allow the TD model to successfully simulate response topography as well as other aspects of NMR conditioning. The chapter also describes simulation results for a number of classical conditioning procedures.

The remaining chapters describe behavioral experiments that bear on predictions of the TD and SBD models. Chapter 4 describes an experiment that evaluated a unique prediction of the SBD model, namely that for fixed-trace interval conditioning longer ISIs can condition more readily than shorter ISIs. The experiments described in Chapters 5 and 6 evaluated a compound conditioning protocol for which the SBD and TD models make different predictions.

Chapter 7 reports general conclusions about the utility of both the TD and SBD models of classical conditioning and describes directions for future research.

CHAPTER 2

ISSUES OF MODELING RESPONSE TOPOGRAPHY

The present research examines the TD and SBD models as descriptions of trial-level features of classical conditioning, such as the development of the CR over trials, and within-trial features of classical conditioning, such as response topography. This chapter first establishes the methodology used in computational work with the TD model and the criteria used for evaluating the performance of both the TD and SBD models. Discussion then turns to the interpretation of V_i (the strength of the association between CS_i and the US) with respect to TD and SBD model output at both trial- and intra-trial levels. The present chapter then focuses on the TD model as a device for producing response topography, showing that in order to generate reasonable response topography, the TD model as proposed by Sutton and Barto (1987) requires changes to the learning rule and to the form of the CS input or representation.

2.1 Methodology

Unless otherwise noted, computational experiments were conducted on the TD model using a simulation program, coded in C, that I developed. The simulator's accuracy was confirmed by replicating the results of Sutton and Barto (1987) (Sut-

ton, 1988, personal communication). In all the experiments reported here, a time step is assumed to represent 10 ms of real time. Such a time frame is consistent with data collection intervals for behavioral and physiological experiments in this laboratory, and is the same as the time frame assumed by Moore et al., (1986) in experiments with the SBD model.

Data derived from the computational investigations reported in this and subsequent chapters include simulated tracings of responses over trials and learning curves (plotted as associative strength over trials). These and other graphical representations of the model's results are depicted in Appendix B.

Model predictions for learning were compared with the literature of the rabbit NMR preparation. Model predictions for response topography were compared with the following empirically observed features of the NMR. The CR rises gradually in a ramped or S-shaped fashion after CS onset to a peak amplitude at or near the time of the US. The response decays rapidly after US offset. By these criteria, simulated responses that are square-shaped or that do not increase smoothly over the ISI are not realistic. Over training, response amplitude increases and response latency decreases. Figure 2.1 illustrates these topographical features with averaged NMRs recorded from a rabbit conditioned with a 650 ms tone CS followed by a 50 ms US. Additionally, at longer ISIs, CRs often have a delayed onset, a feature known as Pavlovian inhibition of delay. There are exceptions to this characterization (see for example Millenson, Kehoe, and Gormezano, 1977) but it is generally applicable to single CS paradigms with ISIs less than 1 sec.

2.2 Mapping V_i onto Behavior

The problem of mapping associative strength V_i onto output is one that has occupied theorists and empiricists for some time (Hull, 1943; Frey and Sears, 1978;

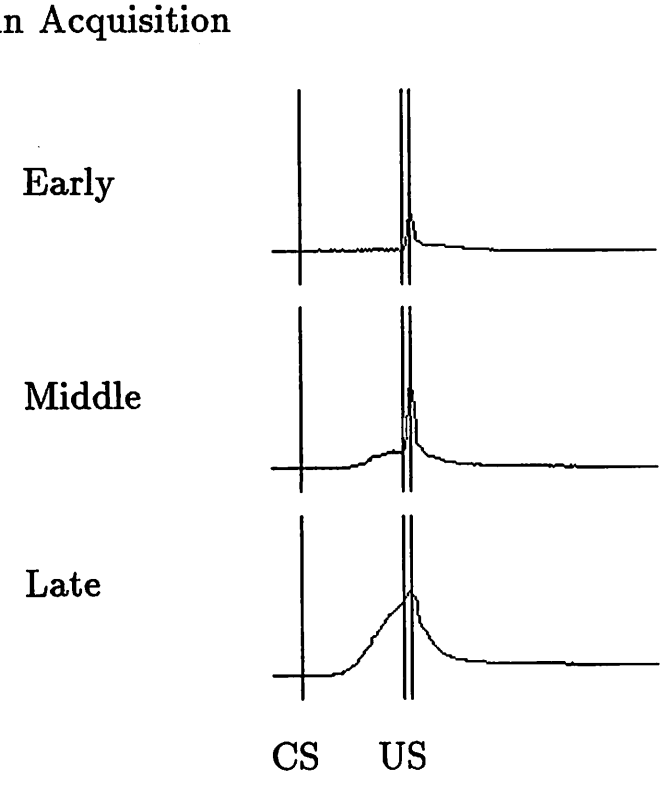

Figure 2.1: **Response Waveforms over Conditioning.** The panels show digitized response waveforms obtained by averaging over 20 trials at different times during training. These data, which are from an animal trained in Experiment 2, were obtained with pairings of a 650 ms CS followed by a 50 ms US at an ISI of 650. Vertical lines show stimulus timing. Early waveform: averaged responses from trials 160-180. Middle Waveform: averaged responses from trials 240-260. Late Waveform: averaged responses from trials from 580-600.

Rescorla, 1988; Tait and Saladin, 1986). For the real-time modeler, there are at least three issues that bear on the mapping of associative strength onto the behavior observed during conditioning. These issues must be considered in light of the response system under observation, since recent studies clearly indicate that the same conditioning procedure can yield excitatory conditioning for some responses and inhibitory conditioning for others (Tait and Saladin, 1986).

The first issue to confront in selecting an output rule is the relationship between V_i and *% CRs obtained over trials*, a typical measure of learning often referred to as a learning curve. Most modelers assume, for the sake of simplicity, that the relationship between acquisition of V_i and the learning curve is one-to-one. However, the empirically observed learning curve is typically sigmoidal, while many models of learning predict that V_i is acquired in a negatively-accelerated fashion.[1] The SB, SBD, and TD models do not produce S-shaped learning curves. Given the ability of these models to account for many other aspects of learning, theorists typically assume that S-shaped learning curves reflect performance rather than learning factors (Moore et al., 1986; see also Frey and Sears, 1978; Rescorla and Wagner, 1972).

The second issue to consider is the comparison of *magnitudes of terminal levels of* V_i obtained with different conditioning protocols. Generally, V_i is taken as an indication of the probability that a response will be evoked by a particular CS. Thus, for example, if V_i obtained for CS1 is greater than that obtained for CS2, we would expect that in a behavioral experiment, the number of CRs elicited by CS1 would be greater than that obtained for CS2. Traditionally, a CS that has a negative value of V_i is thought to be an inhibitory CS, meaning that it may not evoke CRs in the behavioral situation but can be shown to oppose evocation of other CRs in certain behavioral tests (Rescorla, 1969).

[1]Models that do predict S-shaped learning curves include those proposed by Klopf (1988) and Moore and Stickney (1985).

Mapping terminal V_t onto the probability of obtaining a CR may seem reasonable, but does not encompass the occasional simulated conditioning situation for which terminal V_t might be very high but requires thousands of trials to attain. Modelers confronted with situations in which V_t is difficult to interpret may restrict parameter values so as to render V_t interpretable. Narrow ranges of permissible parameter values call into question model robustness but can nevertheless have implications for the kinds of neural machinery required to generate a response (e.g., Moore and Blazis, 1988).

A third issue to consider is *mapping intratrial values of V_t onto response topography*, a problem that is unique to the real-time modeler. To date, most real-time models of learning compute output as the weighted sum of input, including that of the US.[2] Variants of such a rule include those proposed by Schmajuk and Moore (1988, 1989). When a real response system is being modeled, further constraints, such as smoothing and thresholding, are assumed (Moore et al., 1986). In the face of constraints imposed by a particular response system, V_t on a given time step should be interpretable as an index of response amplitude.

2.3 The TD Model and Response Topography

The TD model as stated by Sutton and Barto (1987) is not appropriately configured for the description of intra-trial output that could be interpreted as response amplitude on a given time step. This section presents some of the theory of the TD model and briefly describes the rationale for a modification of its learning rule.

[2]It should be pointed out that another approach to response topography is to map other model variables onto topography. For example, in the TD model, it might be appropriate to map the discrepancy between successive predictions, $\gamma P(v_t, x_t) - P(v_t, x_{t-1})$ (Sutton and Barto, personal communications, 1990). However, the present research assumes that topography is a function of intra-trial V_t.

The TD model is a member of a class of procedures used in prediction problems. It was developed from the Temporal Difference or Sutton's Method (Sutton, 1987).[3] Sutton's method can be applied to a classical conditioning problem in the following way (see Sutton and Barto, 1990, in press, for a full discussion). The TD model assumes that classical conditioning can be viewed as a process of prediction of the US signal, λ.[4] At each time step, the organism predicts the sum of *remaining* values of λ. If time step t occurs halfway through the US, $V(t)$ should be a prediction of the sum of the remaining values of λ. However, since conditioning worsens as ISI lengthens, the TD model assumes that at time t, later values of λ do not influence conditioning as much as more imminent values of λ. Thus, future values of the US signal are discounted. Successive values of λ are discounted exponentially. Therefore, the correct predictions would be of a *sum of the discounted US signals.*

The weight update rule shown in Equation 1.6 and restated here for emphasis provides predictions of the sum of future US values over time, that is, *the future US area.*

$$\Delta V_i(t) = c[\lambda(t) + \gamma P(V(t), x(t)) - P(V(t), x(t-1))]\bar{x}_i(t), \qquad (2.1)$$

where $c > 0$ and $0 < \gamma < 1$. In the simulations, $\lambda(t)$ is 0 when the US is off, and $\lambda(t) > 0$ when the US is on. The CS representation x assumed by the TD model is a rectangular pulse equal to 1 when the CS is on and 0 when the CS is off.

Sutton and Barto (1987) proposed the following rule as for mapping V_i onto model output:

$$s(t) = \sum_{i=1}^{n} V_i(t)x_i(t) + \lambda(t), \qquad (2.2)$$

[3]According to Sutton (1987), the SB and TD models can be generally described as Temporal Difference Methods (TDMs) wherein learning is related to the difference between temporally successive predictions about an outcome. TDMs offer advantages over other methods of prediction: They can be implemented incrementally, thereby conserving memory, and can generalize to new prediction problems more readily than many conventional prediction methods.

[4]The intensity of the US is depicted by the magnitude of λ, and the duration of the US is depicted by the number of time steps in which λ is greater than zero.

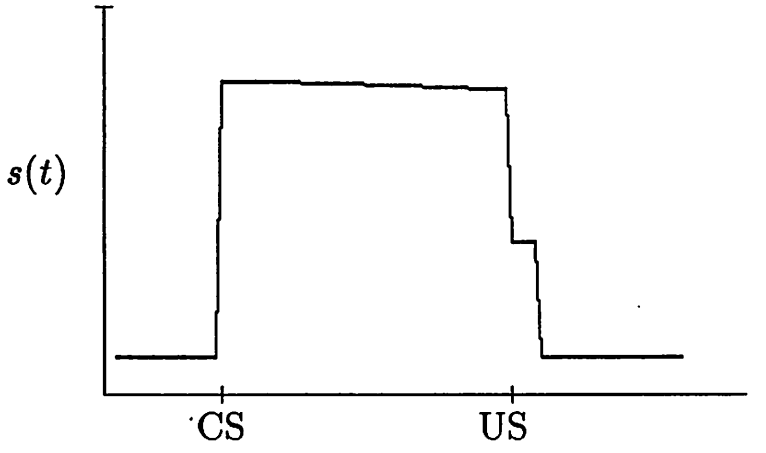

Figure 2.2: **Response Topography with TD Model.** Simulated output $s(t)$ computed by Equation 2.2 during the tenth pairing of a 500 ms CS with a 50 ms US. $\lambda = .3$, $\beta = .95$, $\gamma = .99$, $c = .1$. $V_i = .752$.

Taken together, Equation 1.6 and Equation 2.2 have certain consequences for the generation of response topography. Figure 2.2 shows the response topography generated for a 500 ms CS followed by a 50 ms US, with $\lambda = .3$. Output is computed by Equation 2.2, where x_i is a rectangular pulse equal to 1 when the CS is on and 0 when the CS is off. The figure shows that the CR is substantially larger than the UR. The unrealistically high amplitude of the CR (relative to that of the UR) results because V_i is large, in this case equal to .75 after 10 trials (not asymptotic). V_i is large because the Ps in Equation 2.1 are predictions about the US area, not its magnitude.

Consideration of the large value of V_i led Sutton (1989, personal communication) to suggest that it is appropriate to transform the weights to predict average US signal amplitude instead of US area. Such a transformation was based on the observation that in Equation 2.2, a quantity based on signal area ($\sum V_i(t)x_i(t)$) is being added to a signal amplitude ($\lambda(t)$). The learning rule proposed below in Equation 2.3 transforms the $P(V_t, x_t)$ to predict a discounted average US amplitude instead of US area.

The weight update rule that is consistent with a requirement for an output that can be interpreted as response amplitude is hereafter called the TD_{RT} rule (RT denotes "Response Topography ") and is specified as follows:

$$\Delta V_i(t) = c((1-\gamma)\lambda(t) + \gamma P(V(t), x(t)) - P(V(t), x(t-1)))\bar{x}_i(t), \qquad (2.3)$$

where $0 < c < 1$ and $0 < \gamma \leq 1$. Variable $\bar{x}$ is computed as per Equation 1.8. The derivation of this transformation of V_i is summarized in Appendix C.

One effect of Equation 2.3 is that the terminal value of V_i is always small when typical US durations are used, resulting in topographies that show a marked discontinuity between CR and UR amplitudes (Figure 2.3). The magnitude of V_i does not match the height of the US signal unless the US duration is infinitely long. Empirically, peak CR amplitude is observed to be about equal to peak UR amplitude, even when short duration USs are used. Thus, the prediction of small CR amplitudes shown in Panel A of Figure 2.3 is not desirable.[5]

The assumption of a gain factor greater than 1 alleviates the low CR amplitudes obtained with the TD_{RT} rule:

$$s(t) = \chi \sum_{i=0}^{n} V_i(t)x_i(t) + \lambda(t). \qquad (2.4)$$

where parameter $\chi > 1$.[6] An example of a CR computed by Equation 2.4 is shown in Figure 2.3.

[5]Since normalization leads to low magnitude weights, and also to small amplitude CRs, it might be reasonable to boost the amplitude of the input signal x. However, Chapter 3 shows that increased amplitude of x yields asymptotically lower weights, such that the resultant output is the same height.

[6]Equation 2.4 assumes that output on a given time step is independent of that generated on the preceding step. From a theoretical point of view, it may not be realistic to expect that response amplitude on a given 10 ms time step is independent of the amplitude on a preceding step since there is inertia in the neural and physical systems that generate the response. The assumption of independence in Equation 2.4 may be less problematic when time steps of longer duration are used, thereby allowing the output on a given step to decline fully before the next computation of s.

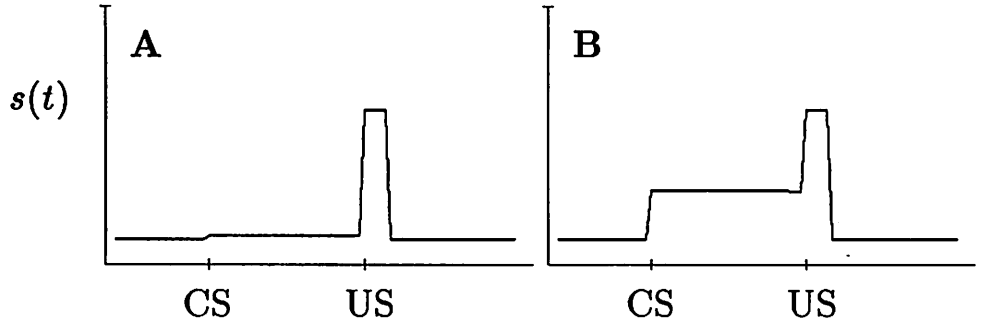

Figure 2.3: **Response Topography with TD_{RT} Rule.** Output $s(t)$ following asymptotic training with a 300 ms CS and a 50 ms US. In A, $s(t)$ is computed with Equation 2.2. In B, $s(t)$ is computed with Equation 2.4, with a gain factor $\chi = 10$. Parameters of Equation 2.3 are set as follows: $\lambda = .5$, $c = .1$, $\gamma = .99$.

2.4 CS Representation and Response Topography

When Equation 2.4 of the TD_{RT} model is applied, output $s(t)$ is a simple function of $x(t)$, the input to the system at time t. Thus, selection of an output rule based on some transformation of the weighted sum of all input creates another issue for the real-time modeler: selection of an appropriate *representation* of CS_i, x_i, that permits the model to properly describe response topography as well as the general outcomes of learning protocols. This section considers two approaches to CS representations: the *rectangular pulse representation* used by Sutton and Barto (1987; 1990, in press) and the so-called *template representation* (Blazis and Moore, 1987; Moore et al., 1986).

As previously mentioned, in implementing the TD model Sutton and Barto (1987; 1990, in press) specified x_i as a rectangular pulse defined by the onset and offset of CS_i. That is, x_i equals one at CS onset and zero when the CS is off. With no further processing, such an input produces an output with zero latency (Figure 2.4).

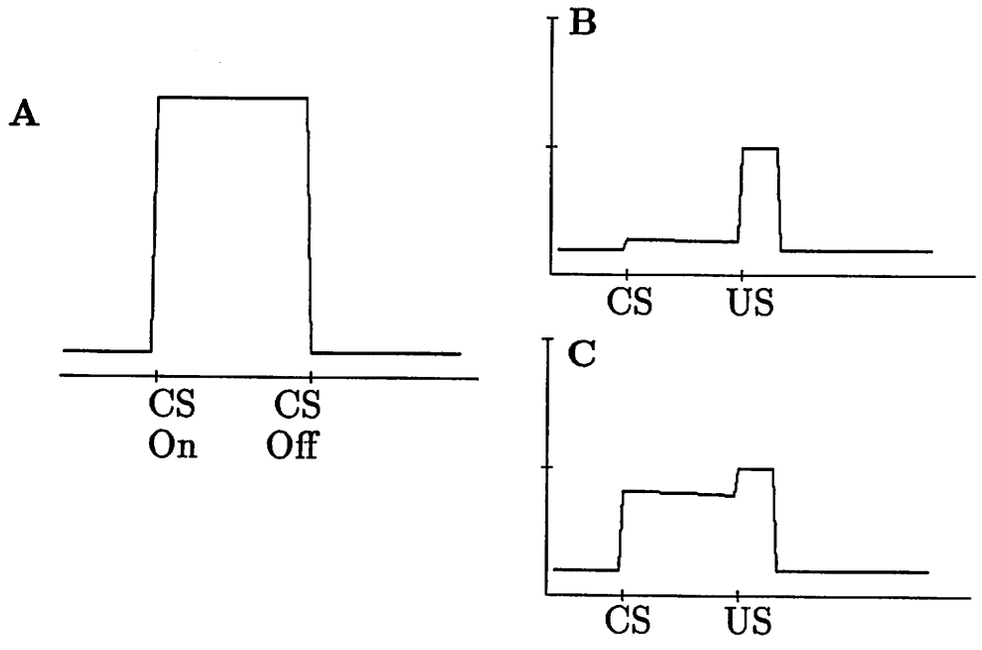

Figure 2.4: **Rectangular-pulse Representation and Development of the CR with the TD_{RT} Model.** Shown in Panel A is the shape of the input of a 300 ms CS, equal to 1 during time steps in which the CS is on and 0 otherwise. Panels B and C show $s(t)$ during the fifth and fifteenth trials, respectively. Tic marks show onset of CS and US; CS ends at US offset. Output $s(t)$ is computed with Equation 2.4, and the gain factor $\chi = 10$.

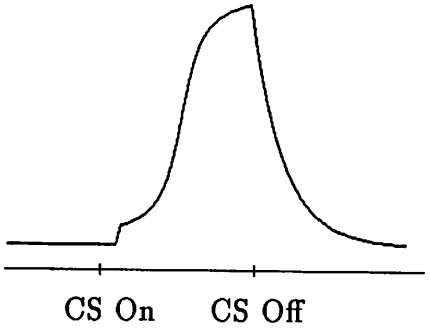

Figure 2.5: **Template CS Representation.** CS representation of the form assumed by Moore et al (1986). Peak amplitude of the input is 1. Tic marks on the x axis indicate CS onset and offset.

A different approach to CS representation is exemplified by the SBD model (Moore et al., 1986). The output rule used in the SBD model is in principle the same as that proposed for the TD model, such that the form of intra-trial output is coupled to the shape of the input. Developers of the SBD model (Moore et al., 1986) sought to overcome the problems of square-wave input noted above by assuming that input to a learning element can serve as a *template* of the response to be learned.

The template for CS input in the SBD model is drawn from the activity of neurons that have been identified in single-unit recording studies as being linked to the generation of CRs. For example, Desmond (1985; see also Desmond and Moore, 1986) investigated the activity of brain stem neurons recorded during classical conditioning of the rabbit NMR with a 350 ms tone CS (ISI of 350 ms). In a typical cell, spikes began to be recruited about 70 ms after CS onset. About 150 ms after CS onset, spike recruitment increased sharply and continued to increase throughout the remainder of the ISI.

Moore et al., (1986) used two functions to model the time course of the CS representation suggested by the physiological data, the first an S-shaped function implemented at CS onset and the second a geometric decay function implemented

at CS offset.[7] The SBD model is able to describe major features of the conditioned rabbit NMR, such as decrease in latency over trials illustrated in Figure 2.1 (see Figure 2 of Blazis and Moore, 1987). The success of the SBD model in this regard is due to the template representation, depicted for a 300 ms CS in Figure 2.5.

In addition to describing learning-related changes in onset latency, topographies generated by the SBD model differ from those generated by the TD and TD_{RT} models. With the SBD model, once peak CR amplitude is reached, CR amplitude is maintained in the time steps prior to the US. However, with the TD model, the CR declines in amplitude in time steps prior to the US.

The potential for loss of V_i during a trial is a feature of the TD and TD_{RT} models in situations where the input level is static. When a CS is on for a period of time, it is eligible for change, and when its input does not increase, the quantity $\gamma P(V_t, x_t) - P(V_t, x_t)$ is negative. By Equation 2.3, high eligibility coupled with a negative difference between these quantities results in a loss of V_i. By Equation 2.4, these decreases in V_i result in loss of CR amplitude. In Figure 2.4, this loss of CR amplitude is evident in Panel C.

The inability of the TD and TD_{RT} models to describe conditioning-related changes in CR latency and the maintenance of CR amplitude over the ISI may be partly alleviated by assuming, as does the SBD model, that the CS representation is an S-shaped function of time. Therefore, the next chapter summarizes simulations of the TD_{RTS} model, a model that incorporates the template (or S-shaped) representation into the TD_{RT} model. In addition to showing that the TD_{RT} model, thusly constrained, predicts response topography, Chapter 3 investigates the possibility that the template representation may adversely affect the TD model's predictions for learning.

[7]The equations defining these functions are considered in the next chapter.

CHAPTER 3

SIMULATION EXPERIMENTS WITH THE TD_{RTS} MODEL

Chapter 2 discussed the TD model, and a variant of the TD model called the TD_{RT} model that produces values of V_i that are normalized for response topography (RT). Like the TD model, the TD_{RT} model cannot describe most features of response topography when the CS is represented as a rectangular pulse. The present chapter summarizes the performance of the TD_{RTS} ("Response Topography with S-shaped Input") model, which is comprised of the following elements:

- TD_{RT} learning rule,
- the template representation x_i mentioned in Chapter 2 and specified below,
- output $s(t)$ computed by Equation 2.4.

Chapter 3 first summarizes simulations studies that examined values for parameters that govern the rise and amplitude of the template representation (described below). Discussion then presents results of simulations studies that investigated the effects of changing parameter values in the learning rule of the TD_{RTS} model. Performance criteria included response topography (as described in Section 2.1), rates of acquisition, and the ISI function, all of which may be affected by the shape of the

template. Finally, the chapter presents the results of simulations studies of a variety of classical conditioning procedures using the TD_{RTS} model. These experiments resulted in further restriction of the set of allowable parameter values.

3.1 A Note About ISI Functions

As indicated previously, one criterion for the performance of real-time models is the ISI function, a ubiquitous feature of classical conditioning. Unless indicated otherwise, the ISI functions examined in this and subsequent sections are for *forward-delay conditioning with the ISI equal to CS duration.* The rabbit NMR literature describes the ISI function as a concave-downward function that rises sharply from the minimal ISI that can support conditioning, approximately 100 ms, to peak at an optimal ISI of 250-300 ms. The function declines gradually as ISI increases beyond the optimal range such that little or no learning occurs at ISIs greater than 2000 ms (Gormezano et al., 1983). Assessed against these criteria, parameter values resulting in ISI functions that are flat or increasing are not allowable.

Given the goal of recapitulating the empirically observed data, one caveat about the experiments discussed below needs to be stated. *In the limit, the TD model, with a rectangular CS presentation, predicts that the ISI function is a decreasing function of ISI* (Sutton, 1990, personal communication).[1] With the template CS representation of the TD_{RTS} model, no learning can occur for CSs of less than 100 ms because x and $\bar{x}$ are nearly zero for the first 100 ms. *Therefore, asymptotically, for ISIs of greater than 100 ms duration, conditioning with the TD_{RTS} model is a decreasing function of ISI,* a prediction not strictly supported by behavioral data. In order to show the ISI function typically obtained with the rabbit NMR, Sutton and

[1] In this dissertation and Sutton and Barto (1987; 1990, in press) data for asymptotic ISI functions with the TD model are not presented. In investigations of the TD_{RTS} model, these data were omitted because in some cases unusually large numbers of training trials were required to reach asymptote.

Barto (1987; 1990, in press) selected parameter values that resulted in a concave-downward ISI function that peaks at an ISI of 250-300 ms.

3.2 Parameters of CS Representation

A battery of simulation studies investigated the performance of the TD_{RTS} model under variations in the parameters that shape the template CS representation described by Moore et al., (1986). The equations specifying the template representation are as follows:

At CS_i onset, $t = 0$. For time steps $t = 1, ..., 7$, the input $x_i(t) = 0$. For $t > 7$ and until CS_i offset, x_i is defined as follows:

$$x_i(t) = [tan^{-1}(mt - 5.5) + 90]/180\, h. \tag{3.1}$$

The parameter m, $m > 0$, controls the rate of rise of x_i, and the parameter h, $h > 0$, controls CR amplitude.

A second function returns the CR generated by x_i to its pretrial baseline. It is implemented at CS_i offset, decays geometrically, and is computed as follows:

$$x_i(t) = kx_i(t-1), \tag{3.2}$$

where $0 < k < 1$.

Simulations summarized in Appendix D investigated the effects of the x-shaping parameters m, h, and k on the TD_{RTS} model's predictions for learning and response topography. The conditioning procedures studied were forward-delay procedures using a 50 ms US and ISI = CS duration. The values selected for simulation were shown by Blazis and Moore (1987) to affect noticeably the SBD model's behavior. Thus, four values of h (0.25, 0.75, 1.0, 2.0), three values of k (0.25, 0.65, 0.85), and three values of m (0.1, 0.35, 0.7) were examined separately and in orthogonal covariations. Remaining parameter values in the TD_{RTS} model simulations were

suggested by pilot work and the simulation experiments of Sutton and Barto (1987) and Sutton and Barto (1990, in press). Specifically, $\gamma = 0.9$, $c = 0.1$, and $\beta = .95$.

As was the case with the SBD model (Blazis and Moore, 1987), the TD_{RTS} model is sensitive to two of the values of parameters that shape the CS representation, m and h. Model performance in the phenomena studied was not greatly affected by the value of k, which determines the decay of the template at CS offset. The following discussion considers each parameter in turn, specifying its optimal value and summarizing the effects of variation from that optimum.

3.2.1 Parameter m

The rise of the template input is governed by the value of m. This parameter was the only one shown to have marked effects on the TD_{RTS} model's predictions for response waveforms and for ISI functions. The simulations presented in Appendix D show that of the three values of m considered, $m = 0.35$ produces the best predictions of response topography and other phenomena.

With $m = .7$, the TD_{RTS} model predicts square-shaped CRs. Little effect on model predictions for ISI functions was noted.

With $m = .1$, CR amplitudes, particularly for CSs of short duration, were especially low. ISI functions obtained with a limited number of trials were unrealistic in that V_i increased with ISI instead of showing the inverted-U shape characteristic of the rabbit NMR. The results obtained with low m were not greatly alleviated by combination with other values of h and k.

3.2.2 Parameter h

The amplitude of the template input is governed by h. The simulations described in Appendix D showed that allowable values for this parameter were $h = 1.0$ or 2.0 (the highest value examined).

With $h = .25$ or .75, CR amplitudes, especially with CSs of longer duration, declined in the time steps prior to the US. ISI functions obtained dropped precipitously such that V_i was very low for the longest CS duration simulated, 800 ms.

The allowable parameter values for the TD_{RTS} model ($m = 0.35$, $h = 1.0$, and $k = 0.85$) also constitute the set of allowable parameter values for the SBD model. This similarity is not surprising because these values generate an S-shaped curve that most closely matches the desired response topography for a short ISI. The fact that x is to be specified in the same way for both models facilitates comparisons between the models.

3.3 Learning Rule Parameters

Having specified the set of x-shaping parameter values most conducive to producing reasonable CR topographies with the TD_{RTS} model, the next set of simulations investigated the effects of varying parameters of Equations 2.3 and 1.8 of the TD_{RTS} model. These parameters include γ, which determines the rate at which future values of λ are discounted, c, the learning rate parameter, and β, the eligibility trace recursion parameter. Acceptable parameter values should permit the TD_{RTS} model to describe ISI functions and reasonable topographies for optimal ISIs. The discussion which follows considers each of the learning rule parameters in turn.

3.3.1 The Discount Rate Parameter, γ

For the TD and TD_{RTS} models, parameter γ determines the amount of increase or decrease in V_i that can occur during the presence of the CS. Its impact can be modulated by the extent to which the CS is eligible and by the size of the learning rate constant c. The resulting fluctuation in V_i during the presence of the CS has consequences not only for model predictions about associative strength over trials, but also for the response waveform within a trial.

Figure 3.1, Panels A and B illustrate the value of V_i over a single trial with the the TD_{RTS} rule and with various values of γ, the discount rate parameter. The effects shown here pertain to both the TD and TD_{RTS} rules, and are relevant for either a rectangular pulse or template representation. Two aspects of the effects of γ are revealed by Figure 3.1, potential for loss of V_i during the CS and the mechanism that generates the ISI function. First, as γ increases, loss of V_i that occurs during the trial decreases. Prior to CS offset, this means that with larger values of γ a CS undergoes less loss of V_i during the trial. At CS offset, the template representation begins its decline. After offset, loss of V_i decreases with increasing γ. Secondly, when CS input levels change slowly, there are more opportunities for unweighting to occur. In Panel A, the CS is 250 ms in duration. In Panel B, a 500 ms CS is shown. The input of the longer CS reaches its asymptotic value of 1 and remains there for about 20 time steps, thereby undergoing more opportunities for unweighting than the 250 ms CS. Therefore, V_i at the end of the trial is lower for the 500 ms CS than for the 250 ms CS at all values of γ.

Figure 3.1 shows no increase in V_i during time steps in which the CS input is increasing sharply; this is because eligibility $\bar{x}$ increases slowly. Figure 3.2 shows the time course of $\bar{x}$ and x for a 250 ms CS. The figure shows that for the template representation considered here, the input indeed increases rapidly during the first

150 ms. By Equation 2.3, this could potentially result in increments of V_i early in the ISI. However, eligibility $\bar{x}$ is low, so potential increases in V_i do not occur.

Asymptotic values of V_i decrease with increasing γ. This is true for the TD_{RT} model when any CS representation is used. With the template representation, for example, when c and β are fixed at .1 and .9, respectively, a 250 ms CS trained with a 50 ms US ($\lambda = 0.7$) has the following weights: .332 when $\gamma = 0.7$, .304 when $\gamma = 0.8$, and .177 when $\gamma = 0.95$. In Equation 2.3, the contribution of the US, $\lambda(t)$, is modulated by the factor (1-γ) in order to produce output appropriate for response topography. As γ increases the contribution of the US decreases, resulting in lower V_i.

Though asymptotic values of V_i decrease with increasing γ for the TD_{RT} model, the same does not hold true for the TD model as specified by Sutton and Barto (1987; 1990, in press). With high γ, the CS undergoes less unweighting during its presence than it does with low γ (see Sutton and Barto, 1987), resulting in larger V_i. Note that the same mechanism operates for the TD_{RTS} model, but that this benefit of high γ is outweighed by the lowered contribution of the US described in the preceding paragraph.

Simulations of the TD_{RTS} model showed that most variations in the value of γ affect the form of the ISI function obtained after 100 trials, and strongly affect response topography (Appendix E). These simulations justified the selection of $\gamma = .95$. With $\gamma < .95$, and assuming the 10 ms time step, ISI functions show greater-than-desired conditioning to short ISIs. Furthermore, with $\gamma < .95$, CR amplitudes decline very rapidly in time steps prior to the US, as discussed above.

3.3.2 The Eligibility Parameter, β

The eligibility trace $\bar{x}$ is computed by Equation 1.8. Parameter β controls the growth and decay of eligibility. Variation in β has some effect on model predictions for the ISI function, but little effect on response topography in the forward-delay procedures investigated by the simulation studies of Appendix E.

Simulations of ISI functions investigated various values of β ranging from .25 to .95. As β increases, V_i for ISIs less than 250 ms decreases, but V_i for ISIs greater than 250 ms increases. This results suggests high values for β, such as .95, are most useful for predicting the rabbit NMR ISI function.

The value of β has implications for ISI functions obtained during trace conditioning with a fixed-duration CS for both the TD and TD_{RTS} models. As β increases, the length of time for which a CS remains eligible following its offset also increases. Thus, the ISI function obtained for trace conditioning is broader with large β (Sutton and Barto, 1987).

Of the values of β examined, $\beta = .95$ was selected as that most conducive to the generation of reasonable predictions by the TD_{RTS} model.

3.3.3 The Learning Rate Constant, c

In the TD_{RTS} model, growth of V_i is determined in part by the value of the learning rate constant c in Equation 2.3. Parameter c influences model predictions for the ISI function and for response topography.

Simulations of the TD_{RTS} model investigated various values of c ranging from .05 to .5. On the basis of simulations studies summarized in Appendix E, values of $c \leq .1$ were considered most appropriate for simulations with the TD_{RTS} model, based on the 10 ms time step. Larger values for c reveal the TD_{RTS} model's tendency towards

a decreasing ISI function after a relatively small number of trials, and produced undesirable declines in CR amplitudes in time steps prior to the US.

3.4 Simulations with the Allowable Parameter Set

Simulation studies investigated the performance of the TD_{RTS} model in a number of conditioning procedures. Parameter values were fixed at those shown in the preceding sections to be prerequisite for reasonable predictions for response topography and ISI functions, namely $\gamma = .95$, $\beta = .95$, and $c = .1$. Unless otherwise noted, the parameter values of the CS representation (Equations 3.1 and 3.2 were set as follows: $m = 0.35$, $h = 1.0$, $k = 0.85$. Unless otherwise noted, for the learning rule, Equation 2.3, the following parameter values were used: $c = 0.1$, $\beta = 0.95$, and $\gamma = 0.95$. Finally, model output was computed according to Equation 2.4.

3.4.1 Acquisition and Extinction of a Single CS

Features of the conditioned NMR for which a model should account include the following: increasing amplitude and decreasing latency of the CR over training, decreasing amplitude and increasing latency of the response during extinction, and the occurrence of peak CR amplitude at or near the US. Figure 3.3 illustrates NMR profiles generated by the TD_{RTS} model over trials for a 300 ms CS. The variable s' on the y axis is the output of the TD_{RTS} model, s, bounded between 0.1 and 1.0. The lower bound of 0.1 imposed on s' represents a threshold due to recruitment effects between the model's output and the motoneurons which generate the peripherally observed response (Moore et al., 1986).[2]

[2]Earlier simulations with the TD_{RTS} model used only s in order to closely examine the model's behavior; the imposition of s' at this juncture facilitates comparison of the SBD and TD_{RTS} models.

At ISIs longer than the 300 ms shown in Figure 3.3, it is possible for CR amplitude to decline in the time steps prior to the US. This decline can be alleviated appreciably by increasing the value of parameter γ.

It is evident from Figure 3.3 that the UR portion of the waveform increases in amplitude over training. The increase in the UR occurs because of x; in this simulation the CS representation is greater than 0 during the US and therefore contributes to output according to Equation 2.4. However, were the US to be presented long after CS offset or by itself, the UR amplitude would remain constant over training. These results differ from those obtained with the SBD model (Moore et al., 1986). The SBD model predicts *diminution* of the UR over conditioning. Conditioned diminution of the UR in the SBD model occurs as a result of the assumption of variable US effectiveness, or λ': US signal efficacy diminishes during trials in which a CS with a non-zero value of V is paired with the US.

The problem of UR modulation, that is, how CSs interact with the US to produce an output during the US, has been a topic of theoretical, behavioral, and physiological research (Donegan, 1981; Donegan and Wagner, 1987; Hupka, Kwaterski, and Moore, 1970; Ison and Leonard, 1970; Weisz and McInerney, 1990; Weisz and Walts, 1990). Either facilitation or diminution of the UR can be observed, depending on parameters of conditioning such as US intensity and ISI. To date, the model that best describes UR facilitation and diminution results is the SOP model of Donegan and Wagner (1987). The models considered in this dissertation provide opposite predictions of UR modulation: The SBD model predicts conditioned diminution of the UR as described above, but the TD and TD_{RTS} models predict conditioned facilitation of the UR. It remains to be seen if both models can be extended to account for both reflex facilitation and diminution.

3.4.2 Interstimulus Interval Effects

Figure 3.4 shows ISI functions obtained with the SBD and TD_{RTS} models. The CS duration was fixed at 300 ms and the US duration was 50 ms. Variable λ = .7. ISIs represented range from –200 ms (backwards conditioning) to 700 ms. The TD_{RTS} model predicts a concave-downward ISI function peaking at about 300 ms. The SBD model predicts the same, but only for ISIs greater than or equal to 200 ms. ISI less than 200 ms result in a negative value of V_i, which is interpreted as inhibitory.

The TD_{RTS} model predicts that, with typical US durations of 50-100 ms, no associative strength accrues to CSs trained at a backwards or a simultaneous ISI. With longer USs (not shown) excitatory conditioning at simultaneous ISIs is possible. Figure 3.4 illustrates that these results are in sharp contrast to those obtained with the SBD model, which, consistent with available data for conditioning of the NMR, predicts inhibitory conditioning at simultaneous and backwards ISIs.

3.4.3 Trace Conditioning

As indicated by Figure 3.4, the TD_{RTS} model accounts for empirical findings of trace conditioning. Conditioning is a concave-downward function of ISI, peaking at 250-300 ms and declining as the ISI increases beyond the optimum. This result is consistent with the literature. However, Chapter 4 presents behavioral data for fixed-trace interval conditioning for which the TD and TD_{RTS} models cannot account. Specifically, for a fixed-trace interval of 300 ms, and using an intense CS, rabbits showed better conditioning at a 1000 ms ISI than at a 650 ms ISI. The TD and TD_{RTS} models predict that V_i declines with ISI (Figure 4.3). Chapter 4 shows that the SBD model, unlike the TD and TD_{RTS} models, provides an account of these results.

Generally, with short ISIs, empirically observed response topographies obtained with trace conditioning tend to begin during or near the end of the CS with peak CR amplitudes occurring in the trace interval. (Chapter 4 shows example topographies). The TD, TD_{RTS}, and SBD models do not predict realistic trace topographies, because the stimulus representation assumed in all of these models provides no explicit representation of the trace interval. Since output (Equation 2.4) is based on a weighted sum of input, and since input is either very low or zero during the trace interval for these models, little or no output can occur during the trace interval.

3.4.4 US Duration and Intensity Effects

As discussed in Chapter 2, the theoretical quantity that is being predicted in the TD model is a discounted sum of the US area. This theoretical basis implies that a low, long US yields as much conditioning as a short, intense US. In spite of these theoretical implications, for the TD_{RTS} model, US amplitude seems to affect conditioning more than US duration. Such a result is consistent with the data (Ashton, Bitgood, and Moore, 1969; Tait, Kehoe, and Gormezano, 1983). Generally, for the TD_{RTS} model, US duration has less of an impact on the increase in V_i because eligibility for increments to V_i declines during the US since eligibility declines at CS offset. Simulation studies and a discussion of this point are shown in Appendix F.

3.4.5 Multiple-CS Conditioning Phenomena

As noted by Sutton and Barto (1987) real-time models can yield interesting predictions for compound conditioning, procedures involving multiple stimuli that occur together but not necessarily simultaneously. Levels of conditioning shown by both the compound and its components are sensitive to component stimulus durations, interstimulus intervals, saliences, and pretraining, as reviewed by Kehoe (1982). Thus the compound conditioning literature provides a rich database for the

assessment of real-time models. A number of compound conditioning procedures involving two CSs were investigated using the TD_{RTS} model.

Conditioned Inhibition

The TD_{RTS} model, like the TD and SBD models, predicts the appearance of Pavlovian conditioned inhibition (CI). Two trial types are involved in Pavlovian CI. One trial type pairs a CS (CSA) with the US; the second trial type pairs CSA with another CS (CSB); the compound is not reinforced (that is, not paired with the US). Over training, the animal learns to respond during the CSA+ trials but not the CSA-CSB trials. CSB has become a conditioned inhibitor; in the model this is represented as having a negative value of V_i. An example simulation is provided in Appendix F. The TD, TD_{RTS}, and SBD models predict that a conditioned inhibitor presented alone will not extinguish, a prediction supported by available behavioral data (Zimmer-Hart and Rescorla, 1974).

Kamin Blocking

In the Kamin blocking paradigm (1968), one CS (CSB, the blocker) is trained to asymptote and then paired with another CS (CSA, for added). The compound of the novel and the pretrained CS is reinforced. Presentations of probe trials are used to assess the associative strength of individual CSs. Generally, very little associative strength accrues to the added CS even when that added CS is more contiguous to the US (Kamin, 1968; Kehoe, Schreurs, and Amodei, 1981). However, if CSA begins before CSB, conditioning to CSA does occur, and CSB undergoes a loss of associative strength (Kehoe, Schreurs, and Graham, 1987). In other words, the temporal primacy of CSA can override the pretraining of CSB. The latter phenomenon has been successfully described by the SB, SBD, and TD models (Kehoe et al., 1987; Sutton and Barto, 1987).

With $\gamma = .95$, the TD_{RTS} model successfully predicts blocking when CSA and CSB occur simultaneously. However, simulations studies showed that in order to account for the temporal primacy effects described by Kehoe et al., (1987), $\gamma = .99$ must be used (see Figure F.4 in Appendix F).

Serial Compound Conditioning

During blocking and conditioned inhibition procedures using simultaneous compounds, the TD_{RTS} model predicts results that are consistent with the literature and does so robustly. However, in certain instances of conditioning with serial compounds, the model requires additional constraints on the value of parameter γ.

An example conditioning procedure in which the TD_{RTS} model's sensitivity is clear is second-order conditioning. In a second-order conditioning experiment, one CS can serve as a reinforcer for another CS. One second-order conditioning procedure involves a two-stage design. In Stage 1 of such an experiment, the first-order CS (CS1) is paired with a US, usually until a reasonably high level of conditioning has occurred. In Stage 2, a new CS (the second-order CS, CS2) is paired with CS1. Responding to CS2 increases initially, but as CS1 extinguishes, responding to both CSs eventually declines (Rashotte, Griffin, and Sisk, 1977). With low γ, it is possible to obtain conditioned inhibition instead of conditioned excitation to CS2. Low γ generates more loss of V_i during CS2 and less increase of V_i during the onset and presence of CS1. With high γ, however, the opposite occurs and CS1 thereby accrues associative strength. An ISI function for second-order conditioning is shown in Appendix F.

3.4.6 Response Topography and the Allowable Parameter Set

Results of the multiple-CS simulation experiments with the TD_{RTS} model dictated that the allowable value of γ is .99. Though as indicated in Section 3.3.1 this increase in γ resulted in lower values of V_i, the topographies that resulted showed no tendency to decline in the time steps prior to the US for forward-delay conditioning with ISI = to CS duration and less than 900 ms. Longer ISIs result in topographies that do sag slightly before the US (see Figure 4.8 in Chapter 4).

3.5 Summary

This chapter summarized aspects of parameter sensitivity in the TD_{RTS} model. Simulations of a variety of single and multiple–CS procedures established allowable parameter values for the TD_{RTS} models. The simulations showed substantial agreement between the TD, TD_{RTS}, and SBD models. The allowable parameter values resulted in reasonable response topographies during acquisition and extinction with the TD_{RTS} model. The TD_{RTS} model tended to show the behavior of the TD model in situations where the latter model's predictions disagree with those of the SBD model.

At asymptote, both the SBD and TD_{RTS} models predict that V_i is a generally a decreasing function of ISI when ISI = CS duration and the CS duration > 100 ms (see Appendix for the data obtained with the SBD model). The asymptotic data were not published in earlier reports (Moore et al., 1986). It should be noted that, with the SBD model, and with a fixed number of trials the SBD model approximates the empirically observed ISI function. The same can be said for both the TD model as presented in Sutton and Barto (1987) and for the TD_{RTS} model with allowable parameter values.

The TD_{RTS} model, like the TD model, contrasts with the SBD model in predicting that simultaneous conditioning is either neutral or slightly excitatory. Most reports from the NMR preparation indicate that simultaneous conditioning is inhibitory but results from other preparations indicate otherwise.

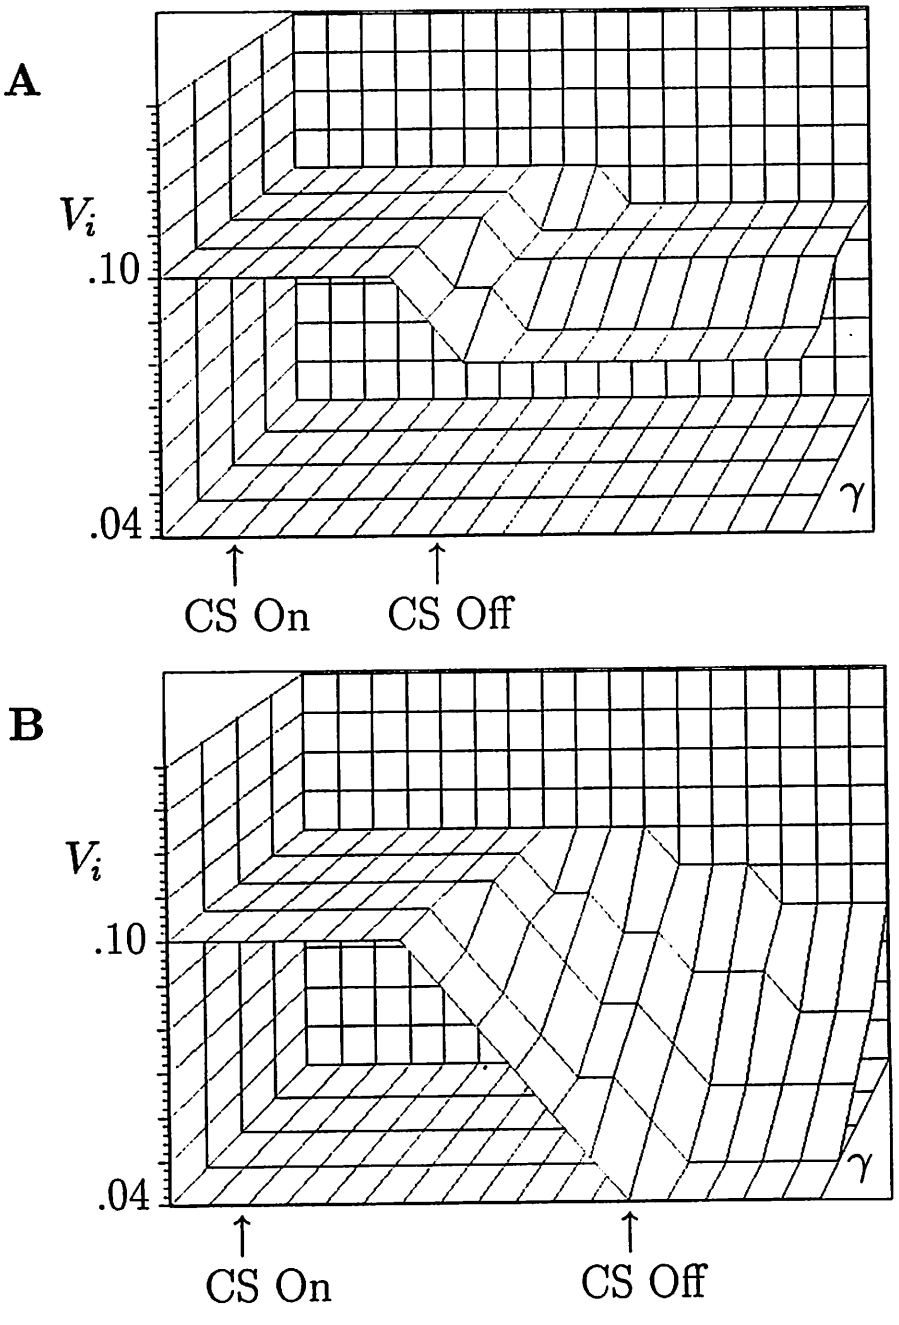

Figure 3.1: **Effect of γ on Intratrial V_i, in the TD$_{RTS}$ model.** Single unreinforced CS; CS input is computed by Equations 3.1 and 3.2. In Panel A, CS duration is 250 ms. In Panel B, CS duration is 500 ms. Shown are the values of V_i (V_i axis) during the trial (horizontal axis) for values of γ (remaining axis, γ = .55, .65, .75, .85 and .95, with .55 closest to reader). Other parameter values included $c = 0.1$ and β = .95.

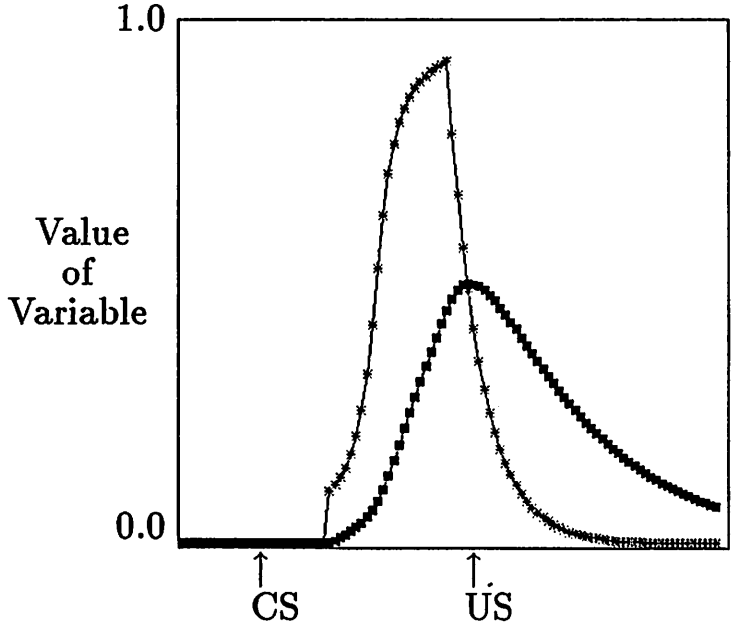

Figure 3.2: **Time Course of x and $\bar{x}$.** The starred plot indicates the time course of x, the CS representation, during a 250 ms CS. Plot denoted with squares shows the time course of $\bar{x}$, the eligibility of the CS. $\beta = .95$.

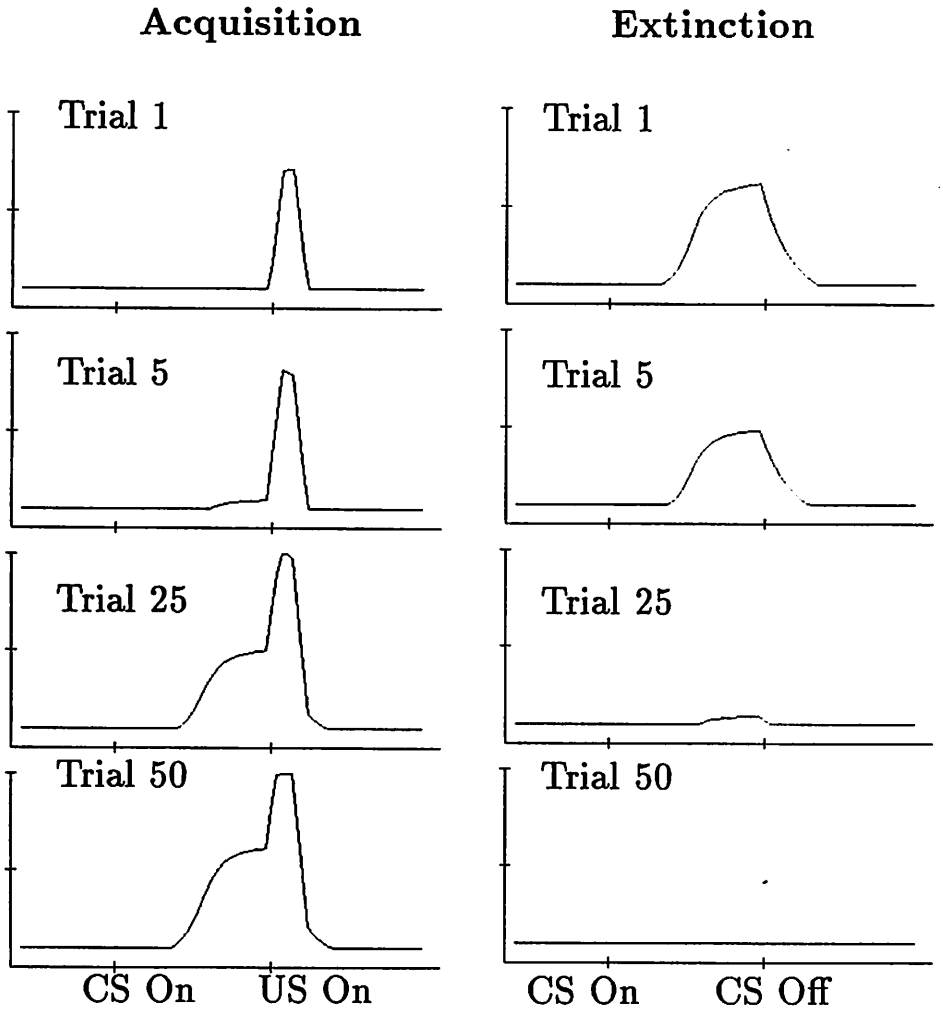

Figure 3.3: **Simulated Response Topographies, s', during Acquisition and Extinction.** These topographies were obtained with the TD_{RTS} model. The CS duration and ISI equal 300 ms. The US duration = 50 ms. The data were generated with the following parameter values: $m = 0.35$, $h = 1.0$, $k = 0.85$, $\lambda = 0.7$; $c = 0.1$, $\chi = 5$, and $\beta = \gamma = 0.95$.

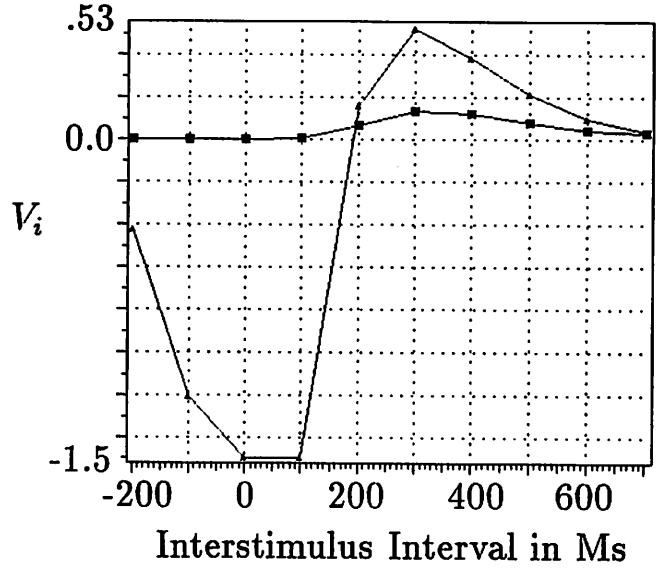

Figure 3.4: **ISI Function for TD_{RTS} and SBD Models with Fixed CS Duration.** Data shown were obtained after 100 trials using a 300 ms CS and a 50 ms US. ISIs range from –200 to 700 ms. Values of V_i generated by the SBD model are indicated by triangles. Values of V_i generated by the TD_{RTS} model are indicated by squares. $\lambda = .7$.

CHAPTER 4

STIMULUS DURATION IN TRACE CONDITIONING

4.1 Overview

Real-time models provide predictions that can be assessed against existing data and detailed recipes for experiments to test novel predictions. As shown in previous chapters, the TD, TD_{RTS}, and SBD models generate the same predictions for many conditioning phenomena, but differences in model predictions also exist. The next three chapters investigate behavioral assessment of procedures for which the SBD and the TD and TD_{RTS} models make differing predictions.

When different models of conditioning are parameterized to predict learning in the same response system, comparison of basic assumptions of the models is facilitated. The preceding chapter showed that the TD_{RTS} model is a variant of the TD model that largely preserves the TD model's behavior as long as a restricted set of parameter values is used. Since the SBD and TD_{RTS} are each specifically constrained for the rabbit NMR, experimental design was based primarily on simulation experiments with these models. In cases where the original TD model offers a different prediction, the difference is indicated.

4.2 Eligibility and Interstimulus Interval Functions

The TD, TD_{RTS}, and SBD models predict fairly realistic ISI functions but do so by different means. For the TD and TD_{RTS} models, appropriate ISI functions result for forward-delay conditioning because CS presence generates loss of V_i (unweighting): the longer a CS remains on, the greater the amount of unweighting incurred, and the lower the terminal V_i obtained. For trace conditioning, the fact that eligibility decreases following CS offset means that for fixed CS durations, lengthening the ISI results in relatively lower values of V_i. For the SBD model, however, appropriate ISI functions result from the assumption that the decay of eligibility following CS offset depends upon the duration of the CS:

$\bar{x}_i$ begins its decline three time steps after CS_i offset:

$$\bar{x}_i(t) = \delta \bar{x}_i(t-1), \tag{4.1}$$

where $\delta = e^{-3/d}$, $d = \max\{25, CS_i \text{ duration in units of 10-ms}\}$. Thus, eligibility at CS offset remains elevated for a longer period of time when a long, rather than a short CS is used.

Figure 4.1 displays the time course of the eligibility trace in the SBD model for forward delay and trace conditioning procedures, using two different CS durations. A near-optimal CS, represented as CS_{350}, is compared against a less-optimal CS, CS_{700}. Also depicted is the value of the discrepancy term $s - \bar{s}$. V_i is computed according to Equation A.3: When the discrepancy $s - \bar{s}$ is positive, V_i increases to the extent that $\bar{x}_i$ is greater than zero. Conversely, V_i decreases when $s - \bar{s}$ is less than zero. Figure 4.1 shows how interactions between the values of $\bar{x}_i$ and $s - \bar{s}$ give rise to particular values of V_i.

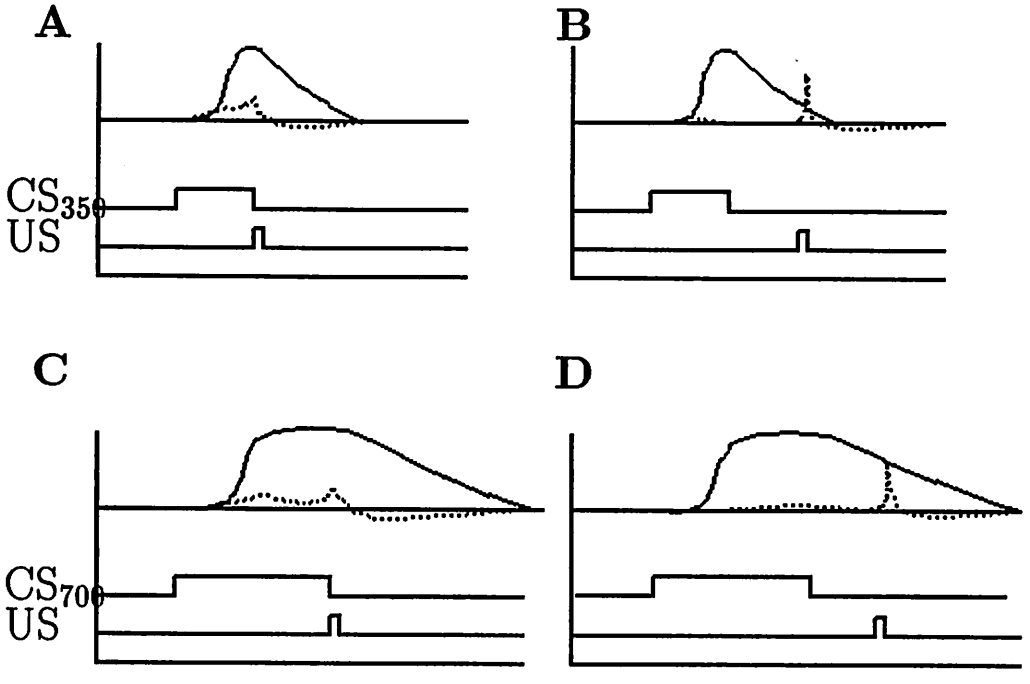

Figure 4.1: **Shape of the Eligibility Trace Assumed by the SBD Model.** The uppermost trace in all four plots represents the form taken by eligibility $\bar{x}_i$ during a trial. The dotted trace represents the value of the discrepancy term $s - \bar{s}$. The topmost straight line represents 0 magnitude for $\bar{x}_i$ and the term $s - \bar{s}$. When $s - \bar{s}$ is greater than 0, V_i increases proportionally to the extent that $\bar{x}_i$ is greater than 0. When $s - \bar{s}$ is less than 0, V_i decreases to the extent that $\bar{x}_i$ is greater than 0. The lower traces show CS and US onsets and offsets. US duration in all plots is 50 ms. In A, the CS duration is 350 ms; the ISI is 350 ms. In B, the CS duration is again 350 ms, but the ISI is 650 ms. In C, the CS duration is 700 ms, the ISI is 700 ms. In D, the CS duration is 700 ms and the ISI is 1000 ms. Values of V_i after conditioning is complete are as follows: panel A, .64; panel B, .12; panel C, .46; panel D, .20.

Figure 4.1, panels A and C demonstrate that $\bar{x}_i$ can interact with the discrepancy term $s - \bar{s}$ to produce results consistent with the empirically observed ISI function: specifically, the weight obtained for CS_{350} (panel A) is greater than that obtained for CS_{700} (panel C). When the decay of $\bar{x}_i$ is extended, as it is with CS_{700}, an extended period of loss of V_i ("unweighting") occurs following the US. Unweighting is due to a relatively high eligibility for modification coupled with the negative value of the term $s - \bar{s}$ during post-US timesteps. The decay of $\bar{x}_i$ is more rapid for CS_{350}, thereby resulting in relatively less unweighting and a higher net value of V_i for that CS.

Figure 4.1, panels B and D shows that the SBD model predicts a procedure in which the empirically described ISI function would *not* hold. That is, a higher value of V_i is obtained for CS_{700} (panel D) than for CS_{350} (panel B). However, the ISI in panel D is longer than that shown in panel B. Since the eligibility trace decays rapidly for S_{350}, very little associative strength accrues at the time of the US, when the value of the discrepancy term $s - \bar{s}$ is high. By contrast, the eligibility of CS_{700} (panel D) is still high due to its slow decay, thus increments in V_i are greater at the time of the US. Substantial decrements of V_i can occur later in the trial, but the net increase in associative strength is greater for the CS_{700} than for CS_{350}.

Figure 4.2 summarizes simulation experiments of ISI effects with the SBD model given four different trace interval durations. When the trace interval is 0 or 100 ms in duration, the SBD model predicts an ISI function that is similar to those empirically observed. However, when a relatively long trace interval of 300 ms is used, the predicted ISI function is shifted, and it is broader; for example, a higher value of V_i is obtained for the 900 ms ISI than for the 450 ms ISI. There are no data concerning this particular manipulation of trace conditioning for the rabbit NMR preparation, and little information is available from other conditioning preparations.

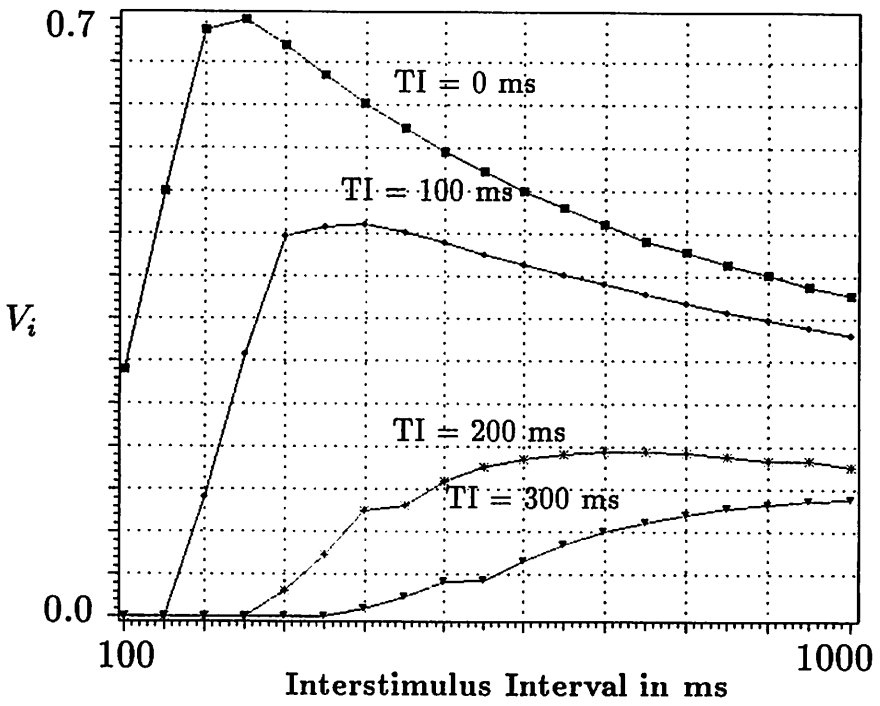

Figure 4.2: **Trace Conditioning with SBD model as a Function of CS Duration.** Magnitude of associative strength V_i obtained after 15 simulated trials in a trace conditioning procedure. Each curve represents the value of V_i obtained for a given trace interval (TI) of 0, 100, 200, or 300 ms duration. CSs range from 100 to 1000 ms in duration. The US is 50 ms in duration and has a λ value of 0.9.

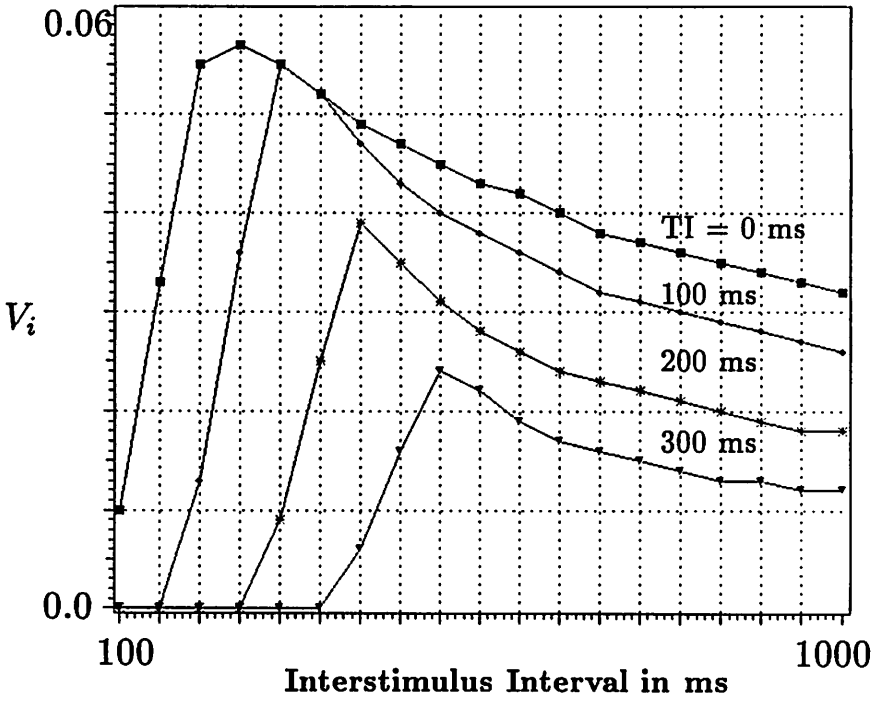

Figure 4.3: **Trace Conditioning with TD_{RTS} Model as a Function of CS Duration.** Magnitude of associative strength V_i obtained after 100 simulated trials in a trace conditioning procedure. Each curve represents the value of V_i obtained for a given trace interval (TI). From uppermost to lowest curve, TI = 0, 100, 200, or 300 ms duration, respectively. CSs range from 100 to 1000 ms in duration. The US is 50 ms in duration and has a λ value of 0.9.

Figure 4.3 shows the results of simulation experiments of ISI effects with the TD_{RTS} model.[1] The figure shows that as the trace interval increases, the peak of the ISI function shifts, but the function itself maintains its inverted-U shape.

If the graphs obtained with the TD_{RTS} and SBD models are compared at TI = 300 ms, it is clear that each model makes a different prediction about the ISI function: For the SBD model, long ISIs condition more readily than short ISIs, but for the TD_{RTS} model, the shorter ISIs condition more readily. Thus, the simulation experiments point to conditioning procedures for which the models yield different predictions, and specify the experimental parameters that should yield the effect behaviorally.

4.3 Experiment 1

4.3.1 Introduction

Experiment 1 tests the possibility, suggested by the simulation studies summarized in Figure 4.2, that for a specified trace interval in trace conditioning, associative strength can increase with CS duration. There are at present no data available to support or disconfirm this prediction for NMR conditioning. Kamin (1965) provides the only study which examines these variables in conditioned suppression. He found no effect of CS duration on trace conditioning in conditioned suppression of ongoing operant behavior, but in his study the trace interval varied so that ISI duration was equalized. Lipkin and Moore (1969) examined the effects of trace interval duration and CS intensity on human eyeblink conditioning; they found evidence for an interaction between trace-interval duration and CS intensity.

[1] In order to compare the TD_{RTS} and SBD model predictions, 100 trials were given in the studies summarized in Figure 4.3. One hundred trials with the TD_{RTS} model and 15 trials with the SBD model yield near-asymptotic conditioning for optimal ISIs.

The simulation experiments summarized in Figure 4.2 suggested optimal experimental parameters. Figure 4.2 shows that when a 300 ms trace interval is used, the model predicts that the level of conditioning achieved with a 700 ms CS will be nearly twice as great as that achieved with a 350 ms CS. Thus Experiment 1 used a 300 ms trace interval, and two CS duration conditions, the 350 and 700 ms conditions.

Since the SBD model does not explicitly incorporate parameters for CS intensity, the model does not imply an effect of CS intensity. Because there may exist an intensity for which the model's predictions do not hold, Experiment 1 included a loud (83 dB) and a soft (63 dB) tone. The literature shows that rate of conditioning increases with CS intensity (Kamin, 1965; Scavio and Gormezano, 1974), but there are no published data pertaining to the effects of CS intensity on NMR topography. Such data were obtained by contrasting the 63dB and 83dB groups in the Experiment 1, and was readily available since the apparatus provided digitized, within-trial topography of the response.

4.3.2 Experiment 1: Method

Subjects

32 naive male and female New Zealand albino rabbits (*Oryctolagus cuniculus*), ranging from 1.5 to 2.0 kg, were obtained from a USDA licensed supplier. Three additional animals were discarded during the course of the experiment due to failure to condition, and one was euthanized following presumed spinal cord injury. Rabbits were housed in cages that exceed USDA regulations and allowed free access to food and water.

Apparatus

Protocols for training of the rabbit NMR were the same as those employed in other laboratories and are patterned after Gormezano's (1966) description. Each subject was restrained in a Plexiglas box and trained individually in 1 of 4 drawers of a fire-proofed, ventilated file cabinet. CSs in this experiment, tones of 1200 Hz, and 63 or 83 dB, were delivered via miniature headphones.

The US was mild electrostimulation of the periocular region, administered across 9mm stainless steel wound clips (Clay-Adams) crimped into the skin inferior and posterior to the eye, and 3mm from the margin. For Experiment 1, the intensity of the US was 1.5 mA(constant current, 50 ms duration, train of .1 ms DC pulses, 10 ms apart).[2]

The control of sequence and timing of stimulus events of the behavioral experiments and data acquisition was achieved through an Apple IIe computer equipped with interfaces and software developed by Scandrett and Gormezano (1980).

Recording of the NMR was achieved via a linear optoelectronic phototransducer which communicates with the eye via a loop of suture sewn into the nictitating membrane; movement of the membrane produces a voltage change that is transmitted to an analog-to-digital converter installed in the Apple IIe computer. Following a standard criterion, a CR was defined as a movement of at least positive .5 mm from baseline any time after CS onset and before US onset.

[2]The electrostimulation US has several advantages. First, parameters of this stimulation are readily controlled. Secondly, the sensory consequences of periocular electrostimulation are well-characterized (Hayashi, Sumino, and Sessle, 1984). Unlike an alternative US, the corneal air puff, electrostimulation does not stimulate the cornea, lashes, or marginal skin. Jerking of the head in response to this US is rarely observed given the current levels of electrostimulation used in our studies. In addition, eyelid retraction, commonly used in connection with air puff US, is not necessary with electrostimulation. Eyelid retraction can contribute to eye irritation, edema, and infection.

Data obtained during the sessions was written to floppy disk and later transferred to a Sun Microsystems local area network. Data were analyzed using software described in Appendix G.

Procedure

All animals received one day of preparation and one day of adaptation to the experimental apparatus. On the preparation day, a small loop of nylon suture (4.0, Ethilon) was sewn into each animal's right nictitating membrane. The hair surrounding the right eye was trimmed and the animal was returned to the colony for recovery. On the preparation day, animals were placed into the conditioning apparatus for a period of 40 minutes (approximately the length of subsequent conditioning sessions). No stimuli were delivered at this time.

Training began on the day following adaptation. Animals were assigned randomly to one of four groups. All groups received trace conditioning. Group L83 was trained with a 700 ms, 83 dB tone followed by the US at an ISI of 1000 ms (trace interval = 300 ms). Group L63 received a 700 ms 63 dB tone followed by the US at an ISI of 1000 ms. Group S83 received a 350 ms, 83 dB tone followed by the US at an ISI of 650 ms (trace interval = 300 ms). Finally, Group S63 received a 350 ms, 63 dB tone followed by the US at an ISI of 650 ms. Animals received 100 trials per day, at an ITI averaging 20 seconds, until a criterion of 90% responding was attained (maximum: 19 days, mean = 13, at least 8 days of training for all subjects).[3]

[3]A possible concern in Experiment 1 would be the inter-trial interval (ITI), an important variable in studies of ISI functions (Levinthal, Tartell, Margolin, and Fishman, 1985). Given a sufficiently long ITI, conditioning of a CS of 2 seconds duration can be as robust as that obtained with shorter, normally optimal durations. The ISI functions on which the models were based models were arrived at via procedures in which the CS was presented 50 to 100 times at ITIs ranging from 25 to 60 seconds (Schneiderman, 1966; Smith, 1968; Smith, Coleman, and Gormezano, 1969). Thus, the ITI for Experiment 1 averaged 30 seconds in duration, and 100 trials per day were given. Experiments which support an ISI by ITI interaction (see Mis, Andrews, and Salafia, 1970) show

4.3.3 Experiment 1: Results and Discussion

Percentage CR and Trials to Criterion

Figure 4.4 depicts the mean percentage of CRs shown by the four groups during the first eight days of training. For groups receiving the 63 dB tone, percentage CRs obtained for group S63 was greater than for group L63, but the difference was not significant (L83, mean = 31.87, SE = 10.10; S83, mean = 53.69, SE = 6.86, t(14) = 1,79, 2-tailed $p < .10$). However, for groups receiving the 83 dB tone, percentage CRs obtained for group L83 was significantly greater than for group S83 (L83, mean = 76.13, SE = 3.15; S83, mean = 54.26, SE = 6.25, t(14) = 3.12, 2-tailed $p < .01$).

Figure 4.5 depicts the mean trials required to attain a number of sensitive criteria: Trials to First CR, Trials to 3 CRs in a row, Trials to 5 CRs in a row, and Trials to 9 CRs out of 10. Figure 4.5 reveals evidence for an intensity dependence of the ISI effect, a characteristic suggested by Figure 4.4 but not borne out by the statistics. For the criteria depicted, the pattern of responding evidenced by animals trained with the loud CS was opposite of that shown by those trained with a soft CS. In all cases, the criteria were reached significantly faster with the *long* ISI when the 83 dB CS was used, and significantly faster with the *short* ISI when the 63 dB CS was used.[4]

---

that the interaction is manifested by ISIs of durations longer than the maximal ISI of 1000 ms used in this experiment.

[4]Results of individual T-tests were as follows. All p values are two-tailed, and comparisons for each criterion were held to a family-wise error rate of .10. For Trials to First CR, Group L83 reached criterion more rapidly than Group S83, $t(14) = 2.922$, $p < .02$. Group S63 reached criterion more rapidly than Group L63, $t(14) = 2.816$, $p < .02$. For Trials to 3 CRs in a Row, Group L83 reached criterion more rapidly than Group S83, $t(14) = 2.921$, $p < .02$. Group S63 reached criterion more rapidly than Group L63, $t(14) = 2.660$, $p < .02$. For Trials to 5 CRs in a Row, Group L83 reached criterion more rapidly than Group S83, $t(14) = 2.447$, $p < .03$. Group S63 reached criterion more rapidly than Group L63, $t(14) = 2.354$, $p < .03$. Lastly, for Trials to 9 CRs in a Row, Group L83 reached criterion more rapidly than Group S83, $t(14) = 2.371$, $p = .03$. Group S63 reached criterion more rapidly than Group L63, $t(14) = 2.244$, $p = .04$.

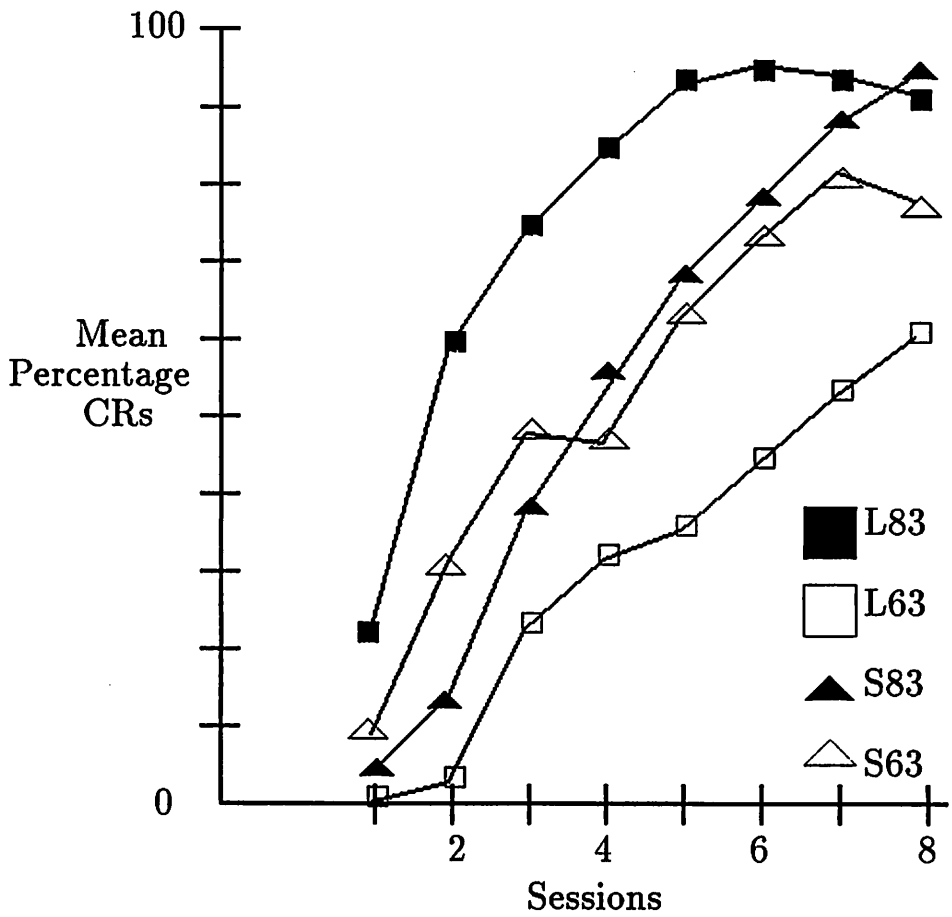

Figure 4.4: **Summary of Experiment 1.** Percentage conditioned responding for each group during each day (Session) of acquisition.

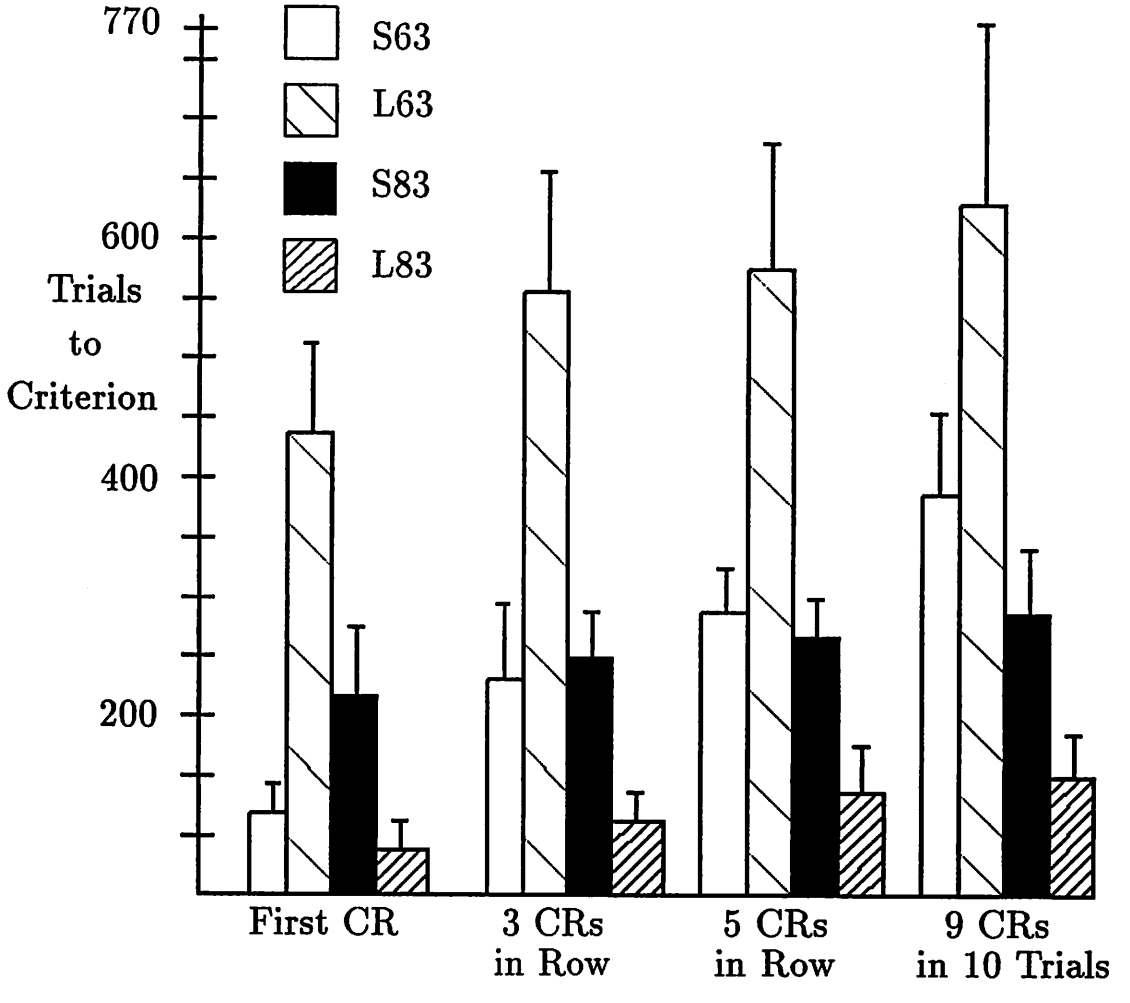

Figure 4.5: **Experiment 1: Trials to Criterion.** On the y axis is shown the mean number of trials required to attain the learning criterion.

Comparison to Model Predictions

Simulation experiments of the SBD model suggest that for certain trace interval durations, it is possible to obtain better conditioning for long rather than short CS-US intervals. The model predicts the behavioral observation of greater mean percentage CRs for group L83 as compared to group S83. However, when a soft CS was used there was no significant difference in the percentage CRs obtained between the long and short ISIs.

The TD_{RTS} model predicts that groups L83 and L63 should have shown worse conditioning (e.g., lower mean percentage CRs) than groups S83 and S63. The present findings do not confirm this prediction.

The results shown in Figure 4.5 show that stimulus intensity and ISI can interact to affect strength of conditioning. Neither the TD_{RTS} or SBD models can account for this interaction. The ISI x Intensity interaction has not previously been demonstrated for the NMR. However, an interaction in trace conditioning between stimulus intensity and ISI duration has been observed using fixed-CS duration trace conditioning in the human eyeblink (Lipkin and Moore, 1966). With a fixed-CS duration and three ISIs (100, 1500, and 2,000 ms) as well as three different CS intensities (70, 80, and 90 dB SPL) Lipkin and Moore found that conditioning was best at long ISIs when a loud CS was used. However, for soft CSs, the shorter ISIs conditioned more readily.

Response Topography

CR Amplitude. Amplitude of NM movement during the ISI in millimeters was recorded for all trials. The average amplitudes after 8 days of training were as follows: Group L83, 2.1 mm, SE = .18; Group S83, 1.6 mm, SE = .21; Group L63, 1.1 mm, SE = .21, and Group S63, 1.3 mm, SE = .15. An analysis of variance

performed on each animal's daily average Peak CR amplitude revealed a significant effect of Days ($F(7,196) = 27.902$, $p < .001$). No other effects were significant, though the main effect of CS Intensity did approach significance ($F(1,28) = 3.939$, $p < .11$).

Response Latency. Response latency was expressed as the time before the US at which the response increased above baseline. After 8 days of training, mean response latency for all eight days was as follows[5]: Group S63, 151.3 ms; Group S83, 119.8 ms; Group L63, 49.8 ms; Group L83, 150.8 ms. An analysis of variance with two between-subject variables (Intensity and Duration) and one within-subject variable (Days) was performed on the average response latency for each of the eight days of training. Consistent with the literature, onset latency decreased over training for all groups. The main effect of Days was significant ($F(7,196) = 47.9$, $p < .001$). The interaction effect of Duration x Intensity was also highly significant ($F(1,28) = 9.57$, $p < .002$).

Peak Latency. Peak latency of a response was defined as the peak magnitude of the NMR waveform during the trial and expressed with respect to US onset. After 8 days of training, the group means were as follows: Group S63, 80.8 ms after US onset; Group S83, 63.9 ms after US onset; Group L63, 50.3 ms after US onset; Group L83, 35.5 ms after US onset. Peak latency decreased over training for all subjects ($F(7,196) = 12.13$, $p < .001$). Animals trained under the long ISI condition showed an earlier peak latency than did those trained under the short ISI condition, as shown by a significant main effect of Duration ($F(1,28) = 34.39$, $p < .001$). Also, the topography of animals trained with the 83 dB CS peaked sooner than did that

[5]Analysis of latency from CS onset is not a meaningful measure in this case because latencies from short ISI groups could never be more than the longest UR latency, which would be in the range of 650-670 ms. However, response latencies for the long ISI groups could be as long as 1000-1030 ms. This would cause an erroneous main effect of duration in analysis of variance.

of the animals trained with the 63 dB CS, as indicated by a significant main effect of Intensity ($F(1,28) = 9.82$, $p < .002$).

CR Area. The CR Area was defined as the summation of A-D counts over the ISI. After 8 days of training, CR area (in units of counts x ms), were as follows: Group S63, 2,515.4; Group S83, 3,424.6; Group L63, 2,361,7; Group L83, 3,412.5. Over sessions, CR Area was found to increase ($F(7,196) = 20.92$, $p < .001$). In general, CR Area increased with intensity (mean for 63 and 83 dB groups, 3,418.5 and 2,438.5, respectively) but this increase was not significant ($F(1,28) = 2.176$, $p < .10$). No other effects approached significance.

UR Amplitude. An analysis of variance was performed on the peak amplitude of the response waveform following US onset. Measures are expressed as mm of NM movement. After 8 days of training group means were as follows: Group S63, 2.9 mm; Group S83, 3.1 mm; Group L6, 2.3 mm; Group L83, 3.1 mm. There was a significant effect of days; as conditioning continued, the peak amplitude increased ($F(1,196) = 10.266$, $p < .001$). Since no probe trials were given, it is not possible to assess the contribution of the CR itself to this effect. No other effects were found to be significant.

4.3.4 Experiment 1: General Discussion

Simulation experiments with the SBD model suggest that in trace conditioning it is possible to obtain better conditioning for long rather than short CS-US intervals. Using the parameters suggested by the model, Experiment 1 tested this prediction in rabbits. The findings indicated that for fixed-trace-intervals, conditioning can be better for long as opposed to short CS-US intervals. The SBD model provides a mechanism for this by assuming that long CS-durations set up traces that can persist over a long trace interval. Other models of conditioning, for example the

TD model, the VET model (Desmond and Moore, 1988; also Desmond, personal communication, 1988) do not make this prediction.

The interaction between stimulus intensity and ISI has not been previously reported for the NMR and is not addressed by the SBD, TD, or TD_{RTS} models, or by any contemporary model. Further explorations of the parameters subserving a CS intensity x ISI interaction are suggested. Experiments employing larger differences between short and long ISIs could indicate the robustness of the interaction effect.

Should the CS intensity x ISI interaction prove to hold over a range of ISIs, it would then be permissible to modify particular variables of the SBD or TD_{RTS} models. One tactic might be to increase c, the learning rate parameter in Equation 1.3, for intense CSs. However, such an action would not change the form of the predicted ISI function, since increases in c have a linear effect on V. Another approach might be to modify the form of variable $\bar{x}$ and its decay to incorporate both intensity and duration parameters.

The variable that would most plausibly encompass CS intensity effects is the CS representation x. Since changes in template representation have undesirable effects on many aspects of the TD_{RTS} model (Chapter 3) and on the SBD model (Blazis and Moore, 1987), alternative CS representations would have to be considered. It should be pointed out that for the TD_{RTS} and TD models, intuitive ways of handling CS intensity can lead to incorrect predictions for V_i. High values of x result in low V_i, while low values of x result in high V_i. This would translate to faster conditioning for less intense CSs than for more intense CSs.

4.4 Verification of the ISI Function

Experiment 1 showed that the shape of the ISI function can change according to the intensity of the CS. Since this is a new result for NMR conditioning, Exper-

iment 2 sought to replicate the empirically observed ISI function for forward-delay conditioning, as reported by a number of laboratories. Though the two ISIs used (650 and 1000 ms) are not very different, the data should show a trend of better conditioning for the shorter ISI in a forward-delay condition. If faster conditioning is observed at the longer ISI, then the results of Experiment 1 may be difficult to interpret in light of existing data.

4.4.1 Experiment 2: Method

Subjects

Sixteen subjects were sutured and habituated to the apparatus as described in Experiment 1. Data from four additional subjects were excluded from the study due to equipment failure.

Apparatus

The apparatus used was the same as in Experiment 1. The US in all cases was the same as that used in Experiment 1, and the CS in all cases was a tone (1200 Hz, 83 dB) delivered via miniature headphones.

Procedure

Preparation and adaptation were conducted as in Experiment 1. The animals were randomly assigned to one of four groups. Group F650 (forward delay, short ISI) was trained with a reinforced 650 ms CS. Group F1000 (forward delay, long ISI) was trained with a 1000 ms CS followed by the US. Two groups were included that replicate Experiment 1: Group T650 received trace conditioning with a 350 ms CS and an ISI of 650 ms, and Group T1000 received trace conditioning with a 700 ms CS and an ISI of 1000 ms (trace interval of 300 ms for both). Animals received

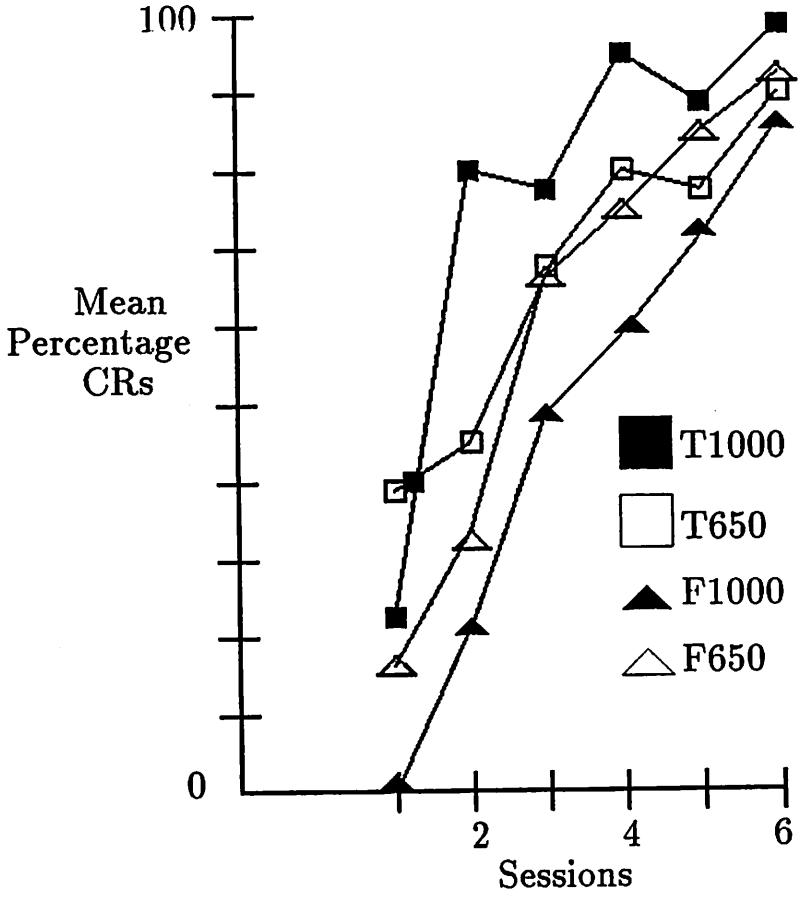

Figure 4.6: Summary of Percentage Conditioned Responding During Experiment 2. Percentage conditioned responding for all groups during Experiment 2; each session is a daily block of 100 trials.

95 reinforced trials and 5 nonreinforced (probe) trials per day, at an average ITI of 20 seconds, until a criterion of 90% responding was attained (mean = 8, maximum = 9, all subjects received at least 6 sessions).

4.4.2 Experiment 2: Results and Discussion

Percentage CRs and Trials to Criterion

Figure 4.6 depicts the percentage CRs obtained for the first six days of conditioning. Mean percentage CRs was as follows: Group F650: 61.5%, SE = 8.0; Group F1000:

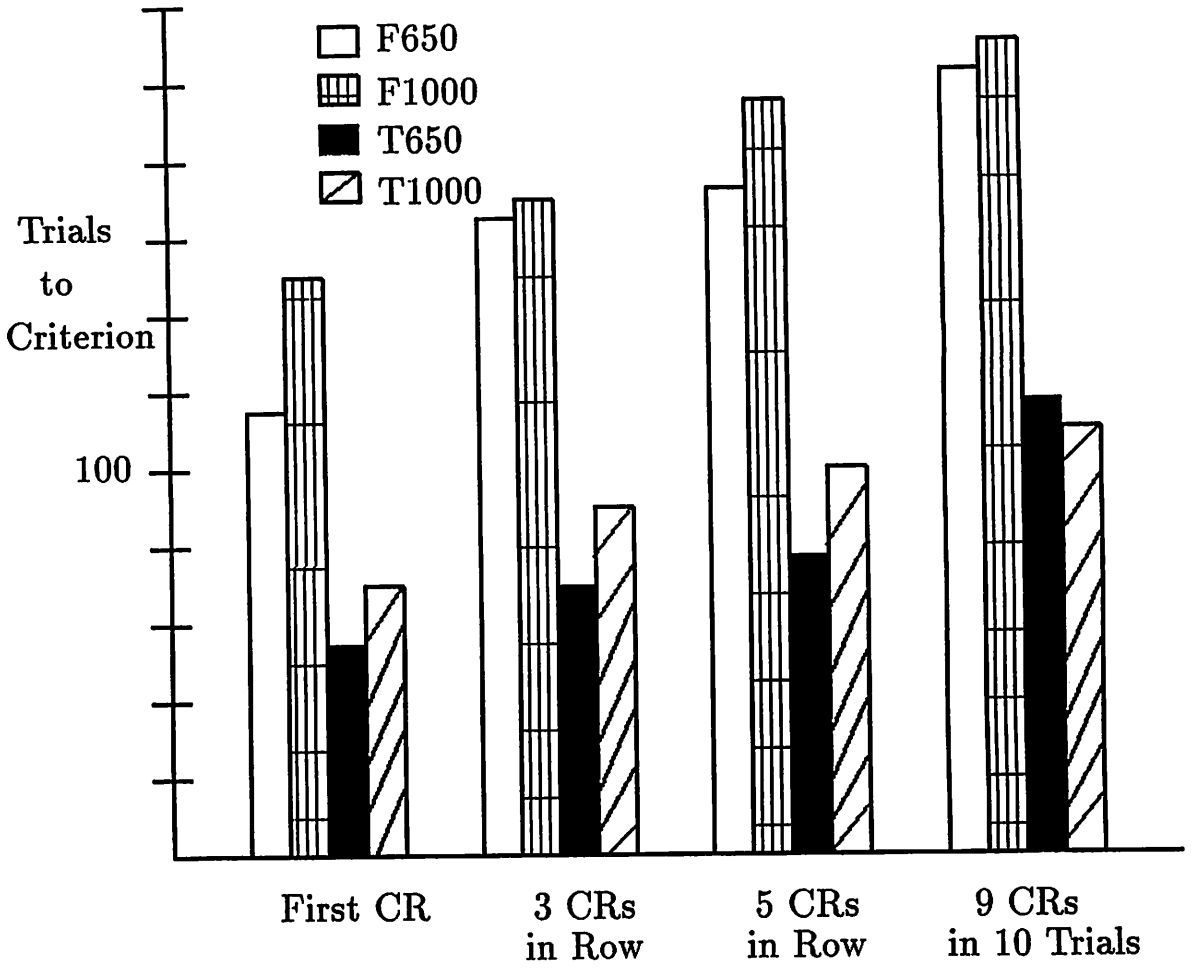

Figure 4.7: **Summary of Trials to Criterion for Experiment 2.** On the y axis is shown the mean number of trials required to attain each learning criterion.

48.2%, SE = 7.0; Group T650, 66.2%, SE = 6.0; Group T1000, 76.8%, SE = 6.0. The figure shows that group F650 conditioned more rapidly than group F1000, and that group T1000 conditioned more rapidly than group T650. An analysis of variance with two between-subject variables (Procedure and ISI) and one within-subject variable was performed on the average percentage CR for the six days of training. The only significant effects were Days ($F(5,60) = 31.16$, $p < .001$) and Procedure (e.g., forward-delay versus trace), ($F(1,12) = 5.71$, $p < .03$). The main effect of ISI was not significant ($F(1,12) < 1$), nor was the Procedure x ISI interaction ($F(1,12) = 2.96$, $p < .10$).

Figure 4.7 shows the mean trials required to attain a number of criteria. For all the criteria shown, trace-conditioned animals required significantly fewer trials to attain criterion than forward-delay conditioned animals[6]. Experiment 2 fails to replicate one result of Experiment 1, that being that for all criteria tested, animals trained under a long-ISI trace procedure learned more rapidly than animals trained under a short-ISI trace procedure, as no significant differences were found between Group T650 and Group T1000. Inspection of Figure 4.7 shows that for most criteria (Trials to First CR, Trials to 3 CRs in a Row, and to 5 CRs in a Row) T650 animals trained more rapidly than T1000 animals. The apparently better performance of Group T650 was due to one subject, T6502, that showed very high responding on the first day that was not continued on subsequent days (see Appendix H). The animal was also unusual in other ways, as he displayed very large and long-duration URs.

The fact that there was no overall *significant* difference between animals trained under the 650 ms versus 1000 ms condition is not surprising given a) the low n used in this study and b) the fact that 650 ms and 1000 ms ISIs are typically only marginally different (Schneiderman and Gormezano, 1964; Schneiderman, 1966).

[6]T-tests, all two-tailed $p < .03$

Response Topography

With the template representation described earlier, neither the TD_{RTS} or SBD models adequately describe the topographies obtained for trace conditioning groups in either this or the previous Experiment. Figure 4.8 shows the topographies predicted by the SBD and TD_{RTS} models after 50 trials (roughly asymptotic). Figure 4.9 shows average response topographies for all subjects on Day 6 of training. Percentage CRs obtained from each group were about the same (Group F650, 90.5%; Group F1000, 86.3%; Group T650, 89.9%; Group T1000, 97%). For all conditions, both models predict that topographies will peak early in the ISI, a prediction not borne out by the data of Figure 4.9. In addition, both models predict that the CR occurs during the CS in the trace conditioning paradigms. The averaged CRs in Figure 4.9 show CRs that occur during the trace interval itself.

CR Onset Latency. The averaged CR waveforms shown in Figure 4.9 do not clearly show features of the onset of the CR in all the groups. The onset of the CR with respect to US onset was examined for CR trials. Group means after 6 days of training were as follows: Group F650, 216.4 ms before US onset; Group F1000, 306.6 ms; Group T650, 147.4 ms; Group T1000, 237.3 ms. An analysis of variance showed that the onset latency for the 1000 ms ISI groups was shorter than for the 650 ms ISI groups ($F(1,12) = 6.780$, $p = .04$). The interaction between Sessions and Duration was significant ($F(5,60) = 4.173$, $p = .006$). Over training, the longer onset latencies of the 1000 ms groups declined precipitously until, on average, onset latencies were much shorter than for short ISI groups.

Trace-conditioned animals tended to respond later in the ISI (Mean for groups T650 and T1000 = 192.4 ms before the US, for groups F650 and F1000, 261.5 ms before the US). Animals trained with long ISIs showed shorter response latencies (Mean for Groups T650 and F650 = 181.9 ms before US onset, Groups T1000 and F1000 = 272 ms before US onset). Consistent with the literature, onset latency

decreased over training for all groups. The TD_{RTS} and SBD models would predict such changes in latency.

Clear pictures of these changes in latency are revealed in Figures 4.10 and 4.11, which plot onset latency distributions for all groups during training days 1, 3, and 5. These distributions show different response tendencies for animals trained with trace or forward-delay arrangements, a trend that is not evident from the analysis of variance. As shown in the upper rows of Figures 4.10 and 4.11, animals trained with a forward-delay arrangement showed a broader distribution of onset latency times, but the distribution tends to peak around the midpoint of the ISI. Animals trained with the 650 ms trace condition showed a well-defined tendency towards responding shortly after CS offset (Figure 4.10, lower row). The same tendency is evident early in training for T1000 animals, but by day 5 the distribution is slightly bimodal, reflecting the tendency of some animals to respond early in the ISI, while others responded later in the ISI (Figure 4.11, lower row).

Peak Latency. As indicated by the predicted topographies in Figure 4.8, the TD_{RTS} and SBD models predict no differences in peak CR latencies. However, the data indicate effects of both ISI and stimulus arrangement on peak CR latency. The mean peak CR latencies for all groups were as follows: Group F650, 52.9 ms before the US, SE = 11.4 ms; Group F1000, 118.1 ms before the US, SE = 18.0; Group T650, 40.2 ms before the US, SE = 6.9 ms; Group T1000, 59.6 ms before the US, SE = 8.6 ms. Subjects trained with forward delay arrangements showed a significantly shorter peak response latency than those trained with a trace arrangement ($F(1,12) = 7.074$, $p < .04$). A significant main effect of ISI indicated that animals trained with the 650 ms ISI had a longer peak response latency than those trained with the 1000 ms ISI ($F(1,12) = 9.997$, $p < .02$). This analysis reveals details of the CR for which neither model can account.

CR Amplitude. Peak CR amplitude was defined as the peak NM movement, in millimeters, recorded during the ISI. These measures were obtained for CR trials; the average peak CR amplitude for each animal was obtained by dividing the sum of the peak amplitudes by the number of CR trials. Average CR Amplitudes were as follows: Group T650, 1.5 mm, SE = .2; Group T1000: 1.5, SE = .2; Group F650, 1.3 mm, SE = .2; Group F1000, .7 mm, SE = .1. An analysis of variance performed on daily average Peak CR amplitude revealed a significant effect of Days ($F(5,60) = 8.465$, $p < .001$). The main effect of arrangement approached significance ($F(1,12) = 3.811$, $p < .15$). No other effects were significant. Figure 4.8 shows that both the SBD and TD_{RTS} models predict differences in CR amplitude according to the group run, but the data do not support this prediction. In fact, the trend in the data would suggest that the amplitudes obtained with the trace conditioning protocols are greater than in the forward-delay conditioning protocols, an effect opposite of that predicted.

4.4.3 Experiment 2: Conclusions

The major finding of Experiment 2 was the replication of the empirically observed ISI function under a loud CS intensity and the stimulus delivery system used in Experiment 1. Two ISIs were used: 650 and 1000 ms. Two conditioning procedures were used: forward-delay and trace conditioning. For the forward-delay condition, animals trained under the 1000 ms ISI condition learned more slowly than animals trained under the 650 ms ISI condition, a finding that is consistent with the literature. For the trace conditioning animals, those trained with the 650 ms ISI trained more slowly than those trained with the 1000 ms ISI, a finding that replicates the result of Experiment 1.

The fact that animals under both trace-conditioned ISIs learned more rapidly than did forward-delay conditioned animals argues against the idea that for most

ISIs, forward-delay conditioning is better than trace conditioning (Schneiderman, 1966). These findings may have resulted from a) the relatively intense CS and/or b) the offset of this salient CS, which occurred at an ISI optimal for NMR conditioning, that is, 300 ms. The finding of better conditioning for trace instead of forward-delay conditioning is consistent with those of Liu and Moore (1969), who showed that the *offset* of a tone 250 ms before US onset was nearly as good a CS as the *onset* of the same tone 250 ms before the US.

Analysis of response topographies showed the shortcomings of the TD_{RTS} and SBD models in predicting appropriate topographies for long forward-delay conditioning and trace conditioning. In all cases, predicted CR onset and peak CR amplitude occurred too early in the ISI, and in the case of trace conditioning predicted CRs occurred during the CS rather than during the trace interval, as was empirically observed. Predictions by the models for differences in CR amplitude were opposite of the CR amplitudes actually observed.

Distributions of onset latencies from the trace-conditioned animals suggest that many of the animals began to respond near CS offset. Though the SBD, TD, and TD_{RTS} models do not currently predict such a result, their CS representations could in principle be modified to generate CRs that occur in the trace interval. An approach that could be useful is that used by Desmond and Moore (1988). The Desmond-Moore model assumes that both CS onset and CS offset are represented as input. This representation enables the Desmond-Moore model to predict appropriately timed CRs during trace conditioning.

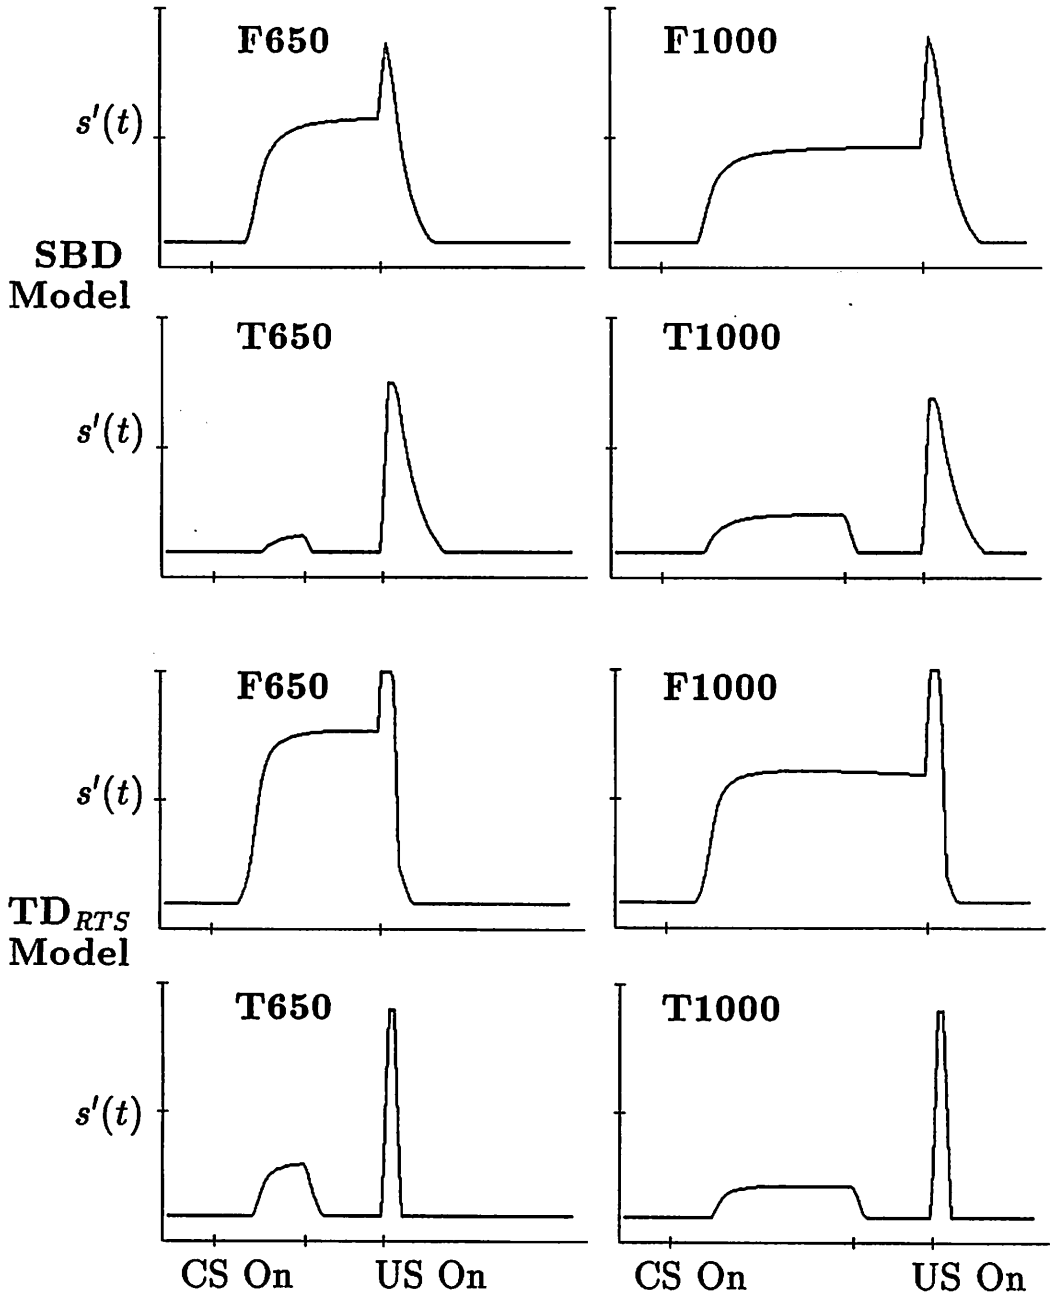

Figure 4.8: **Predicted Response Topographies for Groups of Experiment 2.** Upper plots were obtained with the SBD model after 50 trials. V for each was as follows: F650, .482; F1000, .385; T650, .146; T1000, .207. Lower plots were obtained with TD_{RTS} model using a gain factor χ of 20. V for each group was as follows: F650, .040; F1000, .032; T650, .016; T1000, .012.

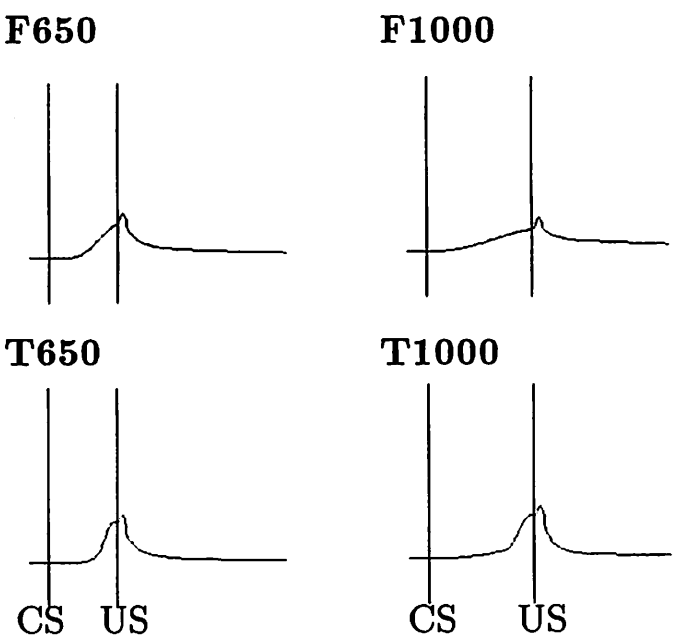

Figure 4.9: **Averaged Response Topographies, Experiment 2, Day 6.** Each panel shows the averaged response waveform for all animals in each group for the sixth day of training. Vertical lines show stimulus onsets. CS offset occurs 300 ms before the US in the lower plots but is not demarcated.

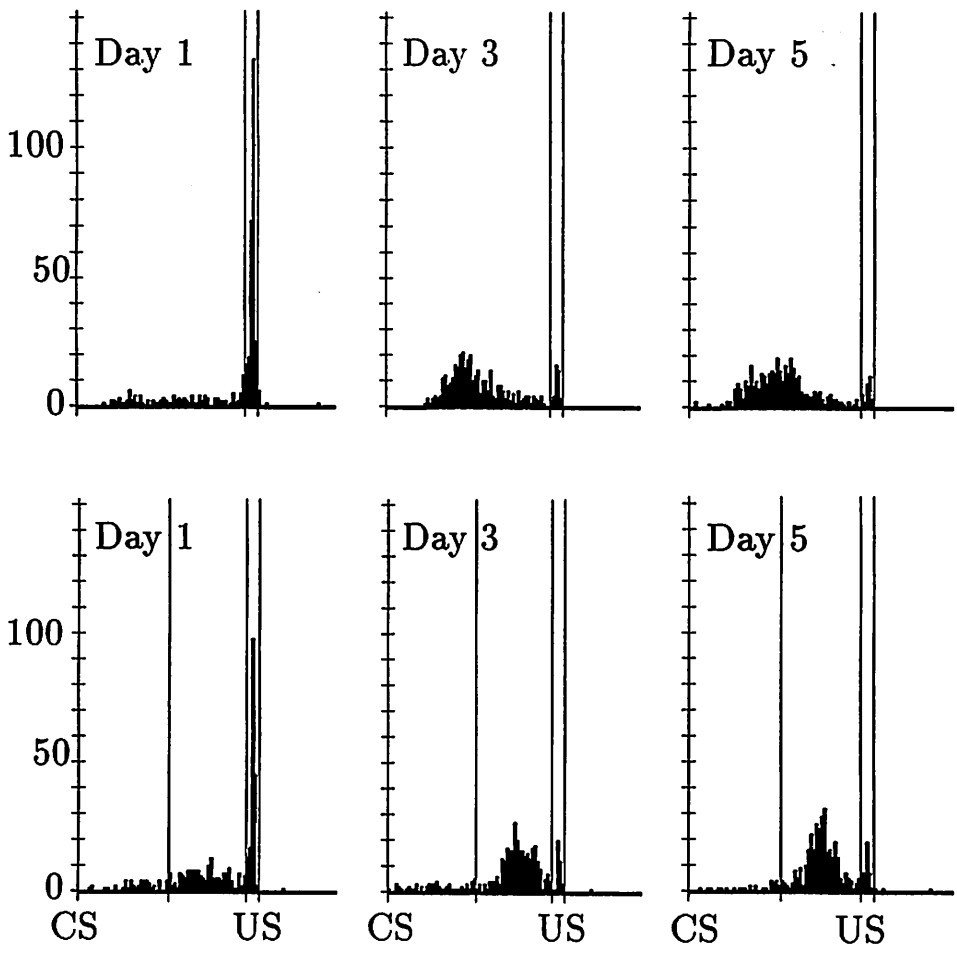

Figure 4.10: **Distributions of Onset Latencies for Experiment 2, 650 ms ISI.** Number of responses at given onset latencies during days 1, 3, and 5 for groups F650 (Upper Panels) and T650 (Lower Panels). CS onset is at the intersection of the axes. Vertical lines show stimulus onsets and offsets. Y axis shows number of responses having a given latency.

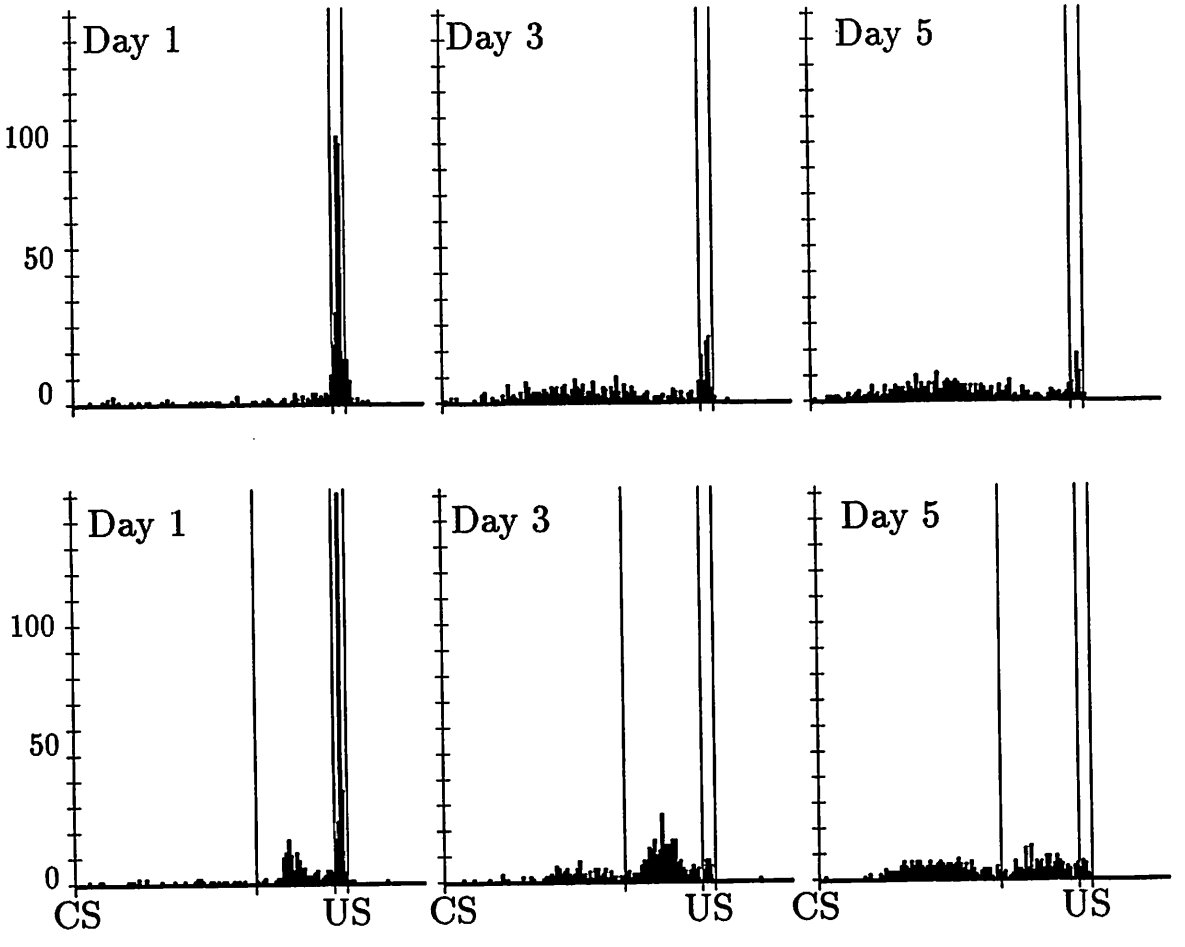

Figure 4.11: **Distributions of Onset Latencies for Experiment 2, 1000 ms ISI.** Onset latencies of all responses during days 1, 3, and 5 for groups F1000 (Upper Panels) and T1000 (Lower Panels). CS onset is at the intersection of the axes. Vertical lines show stimulus onsets and offsets.

CHAPTER 5

NOVEL PREDICTIONS FOR COMPOUND CONDITIONING

5.1 The Complete CS2 Overlap Experiment

As discussed in Chapter 5, compound conditioning involves the pairing of at least two CSs with a US. Learning of the conditioned NMR is very sensitive to the durations of the CSs, the ISIs between the CSs, and the ISI between those CSs and the US (Kehoe, 1982). Therefore, compound conditioning experiments are a rich source of information with which to test real-time models of conditioning.

The present chapter considers a compound conditioning protocol for which SB and TD model predictions are dramatically different. The protocol is here referred to as the Complete CS2 Overlap (CCO) experiment. In the CCO experiment, the first CS, CS1, is followed immediately by both CS2 and the US. CS2 and the US have exactly the same time course.[1]

[1]The CCO experiment is a variant of a simulation experiment, called the Complete Serial Compound (CSC) experiment, conducted by Sutton and Barto (1990, in press). In the CSC experiment, a different CS was presented every single time step before, during, and after the US. Sutton and Barto discussed this experiment because the results are different for the SB and TD models. With the SB model, after many trials the weights of the CSs overlapping the US become negative, effectively cancelling out the US and resulting in extinction of conditioning to the CSs that precede the US. However, with the TD model the weights of the CSs preceding the US increase as a function of proximity to the US; the highest weights are obtained for CSs immediately preceding

Figure 5.1 shows the results of a CCO experiment, run with the SBD,TD, and TD_{RTS} models. Panel A shows the format of the simulation experiment. Two CSs were conditioned; one (CS1) preceded the US and the other (CS2) occured simultaneously with the US. The durations of CS1 and CS2 were 250 and 200 ms, respectively. The US was also 200 ms in duration and with $\lambda = .9$. Figure 5.1, Panel B shows that the SBD model predicts that CS2 accrues negative associative strength. In fact, the negative strength of CS2 can become so great that CS1 extinguishes; in essence CS2 serves to cancel out the reinforcing effect of the US. Panel C shows that the TD model predicts that CS2 accrues positive associative strength. Relative to the single CS control, the TD model predicts greater conditioning for CS1 in the compound condition. In Panel D, the TD_{RT} predicts effects in the same direction as the TD model. However, conditioning to CS1 in the compound is only slightly greater than that which occurs to CS1 in the single-CS condition. In addition to showing that the two models differ in their predictions for the CCO procedure, the simulation experiment shown in Figure 5.1 specifies the training parameters for testing in real rabbits.

In Experiment 3, rabbits were trained using compound conditioning under the CCO procedure. Since the SBD model predicts the appearance of conditioned inhibition to CS2, the experimental design includes summation and retardation tests (Rescorla, 1969). If CS2 is inhibitory, when compounded with a known excitatory stimulus, it should reduce the rate of responding usually observed to the excitatory stimulus (summation test), and, when reinforced, should show slower conditioning in experimental animals relative to control animals that have not previously been exposed to CS2 (retardation test).

---

the US. Weights near the end of and after the US decrease to zero. Sutton and Barto interpret the results obtained with the TD model as a realistic depiction of what might occur behaviorally.

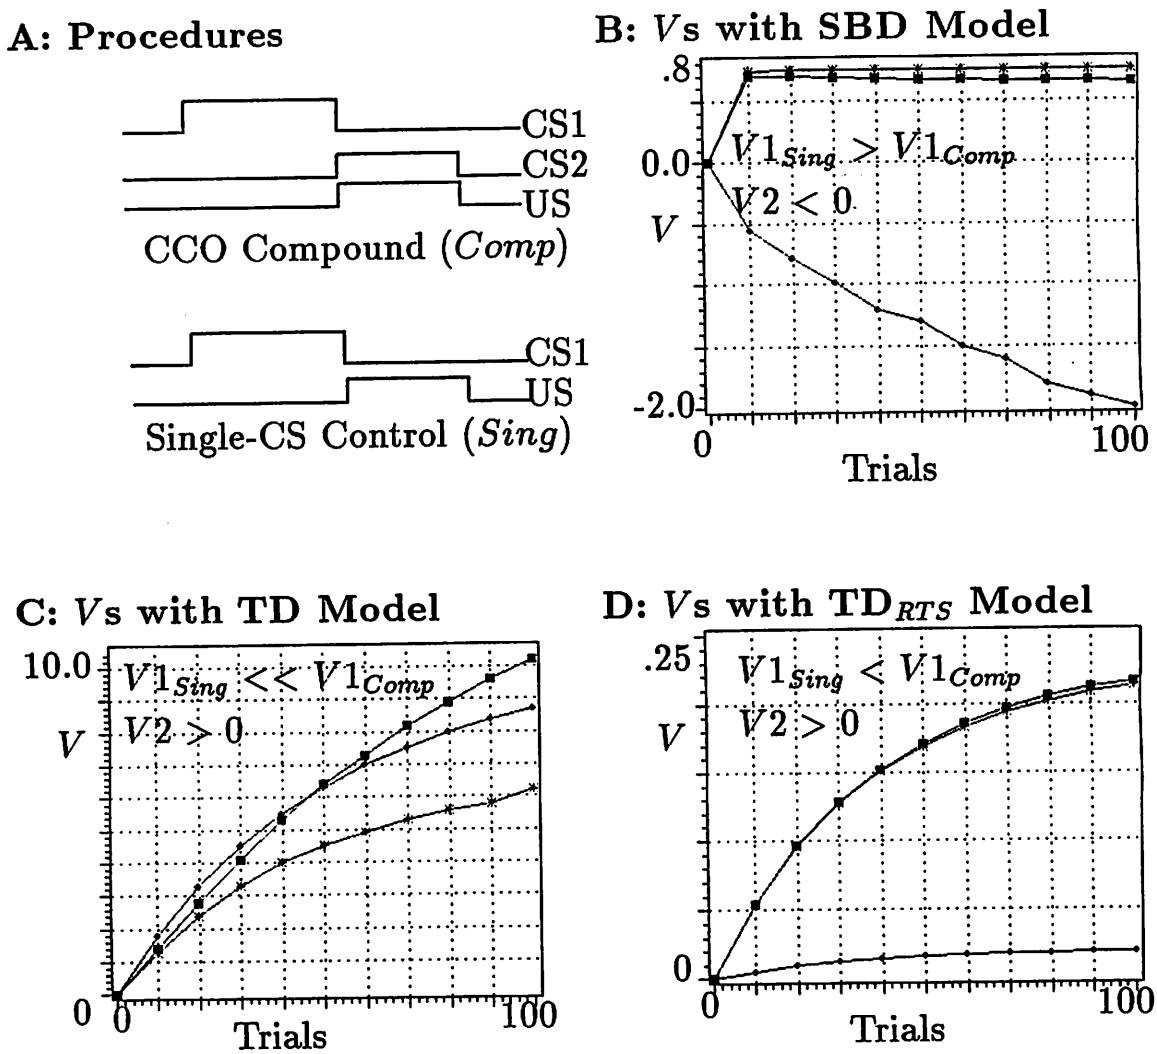

Figure 5.1: **Predictions of the TD,TD_{RTS}, and SBD models for the CCO Experiment.** Panel A shows the timing of CSs and US for the CCO Compound (*Comp*) and the relevant single CS Control (*Sing*). Panels B-D show the predictions for V for each CS in each condition. **Plots with squares** denote $V1_{Comp}$; **starred plots** show $V1_{Sing}$; **plots with diamonds** show $V2$, which is considered only in the Compound condition. In B, simulations with the SBD model. In C, simulations with the TD model. In D, simulations with the TD_{RTS} model.

5.2 Experiment 3

In the the CCO procedure, the first CS in the compound is followed immediately by the second CS and the US; the second CS is coincident with the US. Experiment 3 investigates compound conditioning under the CCO procedure. The stimulus of particular interest is CS2: the TD_{RTS} model predicts that this stimulus becomes excitatory, but SBD predicts that it becomes inhibitory. Since the SBD model predicts the appearance of conditioned inhibition, the experimental design includes summation and retardation tests (Rescorla, 1969).

5.2.1 Experiment 3: Method

16 naive male and female New Zealand albino rabbits were obtained from a USDA licensed supplier. Data from one subject were excluded due to eye infection.

Apparatus

The apparatus used was the same as in the preceding experiments except for the following changes. Acoustic stimuli included a 73 dB white noise CS and a 73 dB, 1200 Hz tone. Instead of using earphones, these acoustic CSs were delivered via speakers located in a panel positioned 10 cm from the subject's head. A light offset CS (designated as "Dark") was also used. The dark stimulus consisted of turning off two 9V incandescent light bulbs located behind milky-white translucent panels, also located 10 cm from the subject's head. The US used in this experiment was the same as in preceding experiments, except that its duration was 200 ms.

Procedure

The procedure used for suturing and adaptation was the same as in the preceding experiments, except that animals were habituated with lights on inside the apparatus (no light was used in preceding experiments). The experimental design is summarized in Table 5.1

Table 5.1. Summary of Design of Experiment 3. Reinforced stimuli are denoted by + signs and nonreinforced stimuli by – signs. Serial compounds are shown by X(Y+) where X (CS1) is the first to occur and Y (CS2) occurs with the US. For all but group T+, $n = 4$. For group T+, $n = 3$.

Group	Stage 1 2 days	Stage 2 12 days	Stage 3 2 days	Stage 4 3 days
DT+	N+	D(T+)	D–, DN, N–,TN – T–, DT–	T+
D+	N+	D+	" "	T+
TD+	N+	T(D+)	" "	D+
T+	N+	T+	" "	D+

Throughout Experiment 3, a CR was defined as a movement of at least positive .5 mm from baseline any time after CS1 onset and before US onset.

Stage 1: Noise Pretraining. The summation test CS, the noise stimulus (CS3), was pretrained for two days. All animals received this training. Animals received 50 trials per day of a 200 ms CS followed by the US. The ITI averaged 45 seconds.

Stage 2: Compound conditioning/controls. Following stage 1, animals were assigned to one of four groups. Two of the groups were compound conditioning groups (corresponding to *Comp* in Panel A of Figure 5.1). Group DT+ received light offset (250 ms) followed by the concurrent presentation of a 200ms tone + 200 ms US at an ISI of 250 ms. Group TD+ received the tone stimulus (250 ms) followed

by light offset + the US. Two single CS controls were also run (corresponding to Group *Sing* in Panel A of Figure 5.1). Groups D+ and T+ received a 250 ms Dark stimulus, or a 250 ms Tone stimulus followed by a 200 ms US at an ISI of 250 ms, respectively. Stage 2 training lasted 12 days. All groups received 5 reinforced presentations of the noise stimulus on each day of Stage 2 training. Hereafter, the noise stimulus is designated CS3, the first stimulus in the compound is designated CS1, and the second stimulus (that which overlaps with the US) is designated CS2.

Stage 3: Summation Testing. On the day following the last day of stage two, all animals received three days of summation testing. Summation testing consisted of nonreinforced presentations of the dark, tone, and noise stimuli as well as all pairwise combinations (e.g., DT, DN, TN). Eight presentations of each type of trial were presented for a total of 48 trials. The CR was scored as a movement of at least positive .5 mm and occurring within 500 ms of CS onset.

Stage 4: Retardation Testing. For the three days following the last day of Stage 3 training, all animals received 50 reinforced presentations of CS2. The ITI averaged 45 ms.

5.2.2 Experiment 3: Results and Discussion

Stage 1: Pretraining

Following Stage 1 mean percentage CRs to the noise CS were as follows: Group TD, 32.1%, Group T, 21.8%, Group DT, 69.5%, Group D, 58%. Due to the exploratory nature of this experiment, no attempt was made to equate the groups for responding to noise. In Stage 2, overall percentage CRs computed for the 12 days were as follows: Group TD, 70.6%, Group T, 75.2%, Group DT, 39%, Group D, 5.2%. Percentage CRs obtained for Group DT was significantly greater than those obtained for Group D (T-test, two-tailed $p < .02$).

Stage 2: Acquisition of Either Compound or Control

Figure 5.2 shows the percentage conditioned responding obtained for each group during Stage 2 training. The acquisition curves of groups TD and T were roughly similiar. However, groups DT and D showed dramatically different patterns of responding: Conditioning proceeded much more rapidly for group DT than for group D. That is, the addition of the tone at the time of the US seemed to boost conditioning to an otherwise poor CS. These results were not anticipated by the SBD model. The TD and TD_{RTS} models predict the facilitation observed in groups DT and D, but these models also predict that faciliation should be observed for groups TD and T as well.

The plots shown in Figure 5.1 showed the predicted growth of V for both CSs over trials. Consider the implications of these plots for Stage 2. CRs were operationally defined as responses at least .5 mm in magnitude before the US, thus the only CRs measured in Stage 2 were CRs to CS1. Since output in both models is computed as a weighted sum of input, and since the input of CS2 prior to the US is zero, the predicted percentage CRs should be a function of the V obtained for CS1. Therefore, the TD model, and to a lesser extent, the TD_{RTS} model, predict greater percentage responding for the compound compared to the single-CS condition. By contrast, the SBD model predicts that a) the control groups should show a higher percentage CRs at all times throughout training and b) the compound groups should show a decline in percentage CRs with continued training. As the data for Stage 2 indicated, the results tend to support the TD and TD_{RTS} models.

Stage 3: Summation Tests

Table 5.2 shows the results of summation tests conducted for all groups in Stage 3. CS1 is the training CS, CS3 is the summation (noise) CS, and CS2 is the stimulus

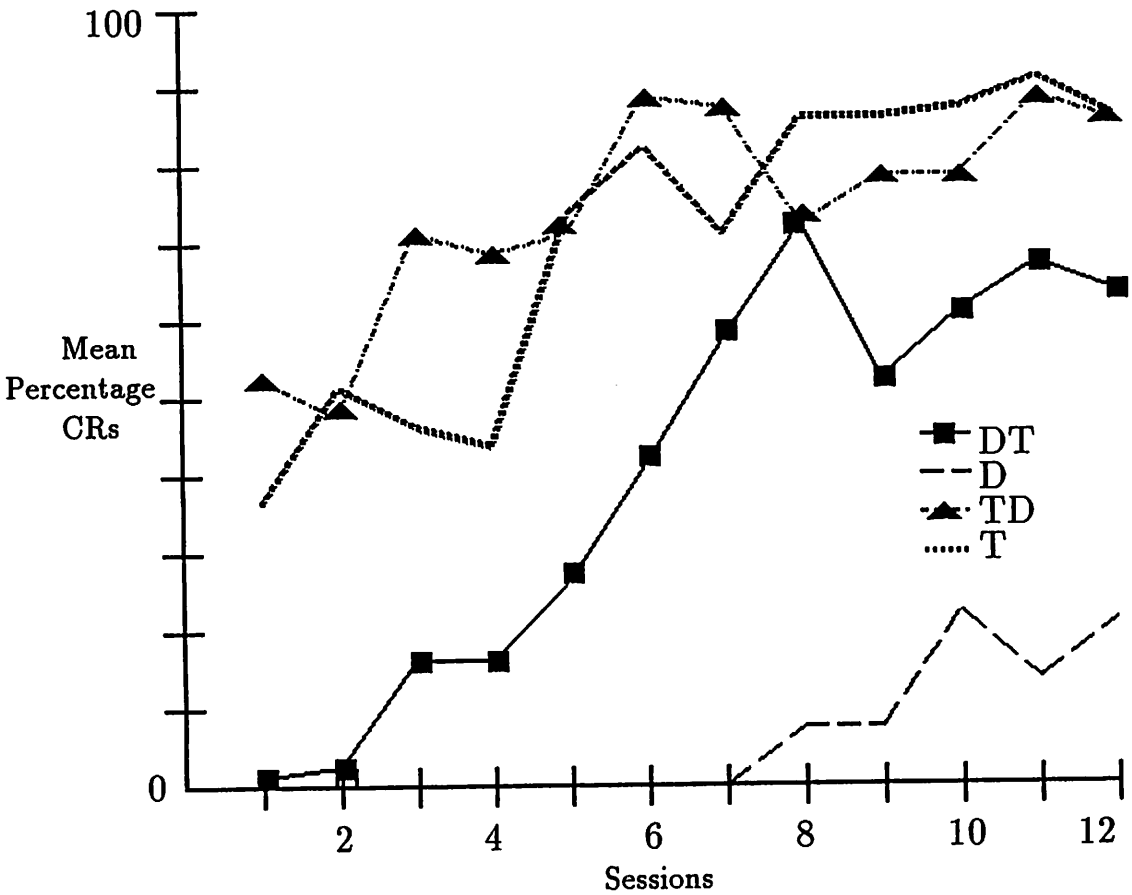

Figure 5.2: **Acquisition of Responding During Stage 2 of CCO Experiment.** Percentage conditioned responding obtained over days during stage 2.

that overlaps the US. Each entry is the group's mean percentage of CRs to the stimulus over three days. As in Figure 5.2, the two single-CS control groups were clearly different from one another. For example, responding to CS1 in the two control groups was markedly different. The dark CS, which served as CS1 in Group D, elicited an average of 11.6% CRs while the tone CS, which served as CS1 in Group T, elicited 61.1% CRs. Since the two groups were different in both Training and Summation (and as shown below, during Retardation as well) the following analysis does not pool the compound and the control groups.

The summation results can be considered in terms of Figure 5.1, which showed the predictions of the TD,TD$_{RTS}$, and SBD models for the CCO experiment. The rank-ordering of the V's can be interpreted as the model's predictions for the summation test:

- **Responding to CS1.** The TD$_{RTS}$ and TD models predict that responding to CS1 under the compound condition should be greater than or about the same as the control condition. The SBD model predicts that responding to CS1 should be lowest for animals trained with the compound.
- **Responding to CS2.** As with CS1, the TD and TD$_{RTS}$ models predict that responding to CS2 in groups trained with the compound should be greater than or about the same as the control groups. This occurs because CS2 accrues at least a small amount of excitatory strength in the compound condition. The SBD model predicts that responding to CS2 should be lowest for animals trained with the compound due to the development of inhibitory strength to CS2.

First, consider the observations for responding to CS1. Referring to Table 5.2, the data show that when Dark was CS1, responding to CS1 was higher for the

compound than for the corresponding control (Mann-Whitney U = 1; two-tailed $p < .05$). As mentioned above in connection with the results of Stage 2, such a result is anticipated by the TD model. For groups TD and T, responding to CS1 (Tone) was slightly greater for group T than for TD, but the difference was not significant.

Observations for responding to CS2 are ambiguous because Dark elicited little responding in Groups TD and T. As for the Tone CS2, responding to CS2 was higher for group DT than for group D, but the difference was not significant.

Table 5.2 Percentage of Conditioned Responding, Summation Testing

Group	Test Stimulus					
	CS1	CS3	CS2	CS1 +CS3	CS3 +CS2	CS1 +CS2
DT	41.7	46.3	24.3	41.6	31.8	33.3
D	11.6	31.4	10.6	31.6	17.6	15.8
TD	43.9	30.2	1.0	41.5	23.9	42.2
T	61.1	56.7	4.2	55.2	59.7	54.2

The SBD model predicts that CS2 becomes a conditioned inhibitor. The important comparisons for establishing CS2 as a conditioned inhibitor are between CS3 alone and (CS1+CS3), and CS3 alone and (CS3+CS2) (Marchant et al., 1972).[2] As a conditioned excitor, CS1 should have increased the likelihood of a CR to CS3. As depicted in Table 5.2, neither the DT or the D group showed excitatory summation, and in fact Group DT showed slight inhibitory summation. As a putative conditioned inhibitor, CS2 should have reduced responding to the CS2+CS3 compound. Both Groups DT and D showed inhibitory summation with respect to level of responding to CS3 (Noise) versus that of (CS3+CS2) (Noise and Tone). The amount of inhibition shown by the DT group is less than that shown by the D group, in spite

[2] A logical non-parametric test for analysing the within-group results of a summation test, the sign test, is not applicable to the present analysis as a minimum of 6 scores are required for a .05 level of significance, two-tailed (Siegel, 1956).

of the fact that for this group responding to both CS3 and CS2 is greater than for controls. This finding suggests that the introduction of a highly salient and novel CS, the Tone, adversely affected responding to the (CS3+CS2) compound in the Group D.

The results obtained for groups TD and T were consistent with the SBD model's prediction of development of inhibitory associative strengths for CS2. Again, the important comparisons are between CS3 alone and (CS1+CS3), and CS3 alone and (CS3+CS2). Group TD showed excitatory summation with respect to CS3 (30.2%) versus (CS1+CS3) (41.5%). Responding to CS3 and the (CS1+CS3) compound was about the same for Group T. The T group showed slight excitatory summation with respect to CS3 versus the compound CS3+CS2 (56.7% versus 59.7%). However, the TD group showed inhibitory summation with respect to CS3; responding to CS3 was 30.2%; to the compound (CS2+CS3) 23.9%; the decline in responding from CS3 to compound was 21.9%. The control group, T, showed a slight increase in responding from CS3 to the compound of 5.3%. This trend, however weak, is consistent with the development of conditioned inhibition and is predicted by the SBD model.

Retardation Test

The results of Experiment 3 suggest that CS2 was a conditioned inhibitor, but the results of the retardation tests are ambiguous. For groups TD and T, mean responding to the Dark CS2 did not exceed 5% during the three days allotted for retardation testing. Table 5.3 shows the average percentage conditioned responding during T+ retardation training for groups DT and D. Group D, the single-CS control, showed a greater mean percentage responding over all three days of training, in spite of the fact that responding to T probes during summation was lower for Group DT (10.6% than for group D, 24.3% for group DT). Such a result suggests two things: a) The inhibitory summation with respect to CS3 shown by Group D

was probably an attentional effect, meaning that the CS2 distracted attention from CS3 in the compound; b) The inhibitory summation shown by Group DT with respect to CS3 was probably *not* a result of distraction of attention by CS2 because Group TD conditioned more slowly during retardation. Therefore, it is possible that for Group DT, CS2 was a conditioned inhibitor.

Table 5.3. Retardation Test for Groups DT and D.

Group	Day 1	Day 2	Day 3
DT	49.2	78.5	74.2
D	58.0	79.7	87.9

Given the failure of the Dark CS2 to condition during retardation testing, a second replication of this experiment was not performed.

5.2.3 Experiment 3: Conclusions

Experiment 3 was designed to assess the differential predictions of the TD, TD_{RTS}, and SBD models in the CCO experiment, a serial-compound training procedure for which one CS had the same time course as the US.

Experiment 3 showed facilitation of responding to a compound containing a weak CS1+strong CS2 relative to the single-CS control. For animals that received a tonal CS1, there was no difference in acquisition relative to a single-CS control. However, for animals that received the Dark CS as CS1, addition of the tone at the time of the US enhanced conditioning relative to single-CS controls. Paradoxically, summation tests for these animals suggest that CS2 might have been a conditioned inhibitor, and such a result is consistent with the SBD model.

However, a number of factors preclude a clear statement as to whether CS2 was excitatory or inhibitory. The unexpectedly low salience of the dark CS precluded

timely retardation tests for groups TD and T, in spite of qualitative evidence that CS2 had passed the summation test for a conditioned inhibitor. The possibility of generalization of conditioning from the noise to the tone, suggested by very rapid acquisition of T+ by groups DT and D, presents another obstacle to assessing inhibition for these groups.

The TD model, and to some extent, the TD_{RTS} model, anticipates the enhanced conditionability of the DT compound relative to the D control (Section 5). Two caveats should be mentioned. The first caveat concerns the TD_{RTS} model. With the TD_{RTS} model and for CS1, the mapping of the model-generated learning curves onto the empirically-observed learning curves is not completely accurate because the model-generated learning curves are the same until late in training (Figure 5.1). The predicted enhancement of conditioning to CS1 late in training occurs because the onset of CS2, a stimulus which actually accrues slight associative strength, adds to the pool of positive reinforcement available for CS1. The second caveat is that both the TD and TD_{RTS} models predict facilitation for the compound relative to the single CS control regardless of saliences of CS1 and CS2. In Experiment 3 the facilitation of responding to CS1 was restricted to the compound in which CS1 was less salient, that is, Group DT.[3]

Experiment 3 thus provides support for the TD and TD_{RTS} models in that facilitation of responding can be seen in the CCO compound relative to its single-CS control. However, some support for the SBD model is suggested by the weak evidence that the overlapping CS was inhibitory.

[3]As an example of the fact that salience changes do not affect the TD model's prediction, consider the following simulation experiment in which salience was presumed to increase the learning rate. For the TD model, if CS1 is twice as salient as CS2 (learning rate for CS1 is .2; that for CS2 is .1) then V's for the procedure illustrated in Figure 5.1 are as follows: $V1_{Comp} = 13.9$, $V2 = 7.9$, $V1_{Sing} = 10.2$. If CS1 is one-half as salient as CS2 then V's are as follows: $V1_{Comp} = 13.4$, $V2 = 8.5$, $V1_{Sing} = 10.3$. The magnitude of the difference in $V1$ between the groups is acutally greater when CS1 is more salient than CS2, due to increasing the potential loss of V that occurs during CS2 presence.

CHAPTER 6

THE CCO EXPERIMENT: FACILITATION REVISITED

6.1 Experiment 4

In Experiment 3, compound conditioning under the CCO protocol, using a weak CS (Dark) followed by coincident presentation of a strong CS (Tone) and a US, yielded significantly greater conditioning than the corrresponding single CS control (Dark). The results of Experiment 3 were not fully anticipated by the TD, TD_{RTS}, or SBD models. Therefore, Experiment 4 further investigated the observed facilitation.

Experiment 4 was designed to alleviate features of Experiment 3 that complicated interpretation of the CCO experiment. In Experiment 3, animals were first pre-trained to a noise stimulus. Thus, the potentiation observed in group DT (Dark+Tone) may have reflected some type of generalization between the noise-pretraining and the tone appearing in the DT compound. Experiment 4 investigated the possibility that faciliated conditioning to the compound could occur in the absence of pretraining. Since Experiment 3 showed that the dark CS is a very-low salience stimulus, Experiment 4 used a relatively more salient light CS as CS1.

6.1.1 Experiment 4: Method

Subjects

Subjects were twenty-four rabbits weighing between 2.0 and 3.1 kg.

Apparatus

The apparatus was the same as in Experiment 1. The US was the same as that used in Experiment 3 (200 ms, 1.5 mA DC electrostimulation of the area around the eye). The CSs used were a tone (1200 Hz, 73 dB) and light, produced by a single 6-Volt light bulb located to the right front of the subject's head.

Procedure

Preparation of the subjects and adaptation to the apparatus were conducted as in preceding experiments. Adaptation was conducted with light off, unlike Experiment 3. Prior to adaptation, animals were assigned so as to roughly match weight and temperament (as ascertained during insertion of the suture). The animals were assigned to one of three groups. Group PLT, the Experimental group, received a 250 ms light (L) followed by the coincident presentation of a 200 ms tone (T) and a 200 ms US (ISI between L and T = 250 ms). No stimuli were presented during the ITI, but data were collected at that time. Group PLP, a single-CS control, received a 250 ms light (L) that immediately preceded (P) a 200 ms US; the ISI was 250 ms. Finally, Group PLO, another single-CS control, received a 450 ms light (L) that overlapped (O) and coterminated with the 200 ms US at an ISI of 250 ms. For Groups PLP and PLO, a 200 ms tone was presented during the intertrial interval. This was done to ensure equal exposure to the tone over all groups. The ITI averaged 30 seconds. Animals received 50 trials per day for seven days.

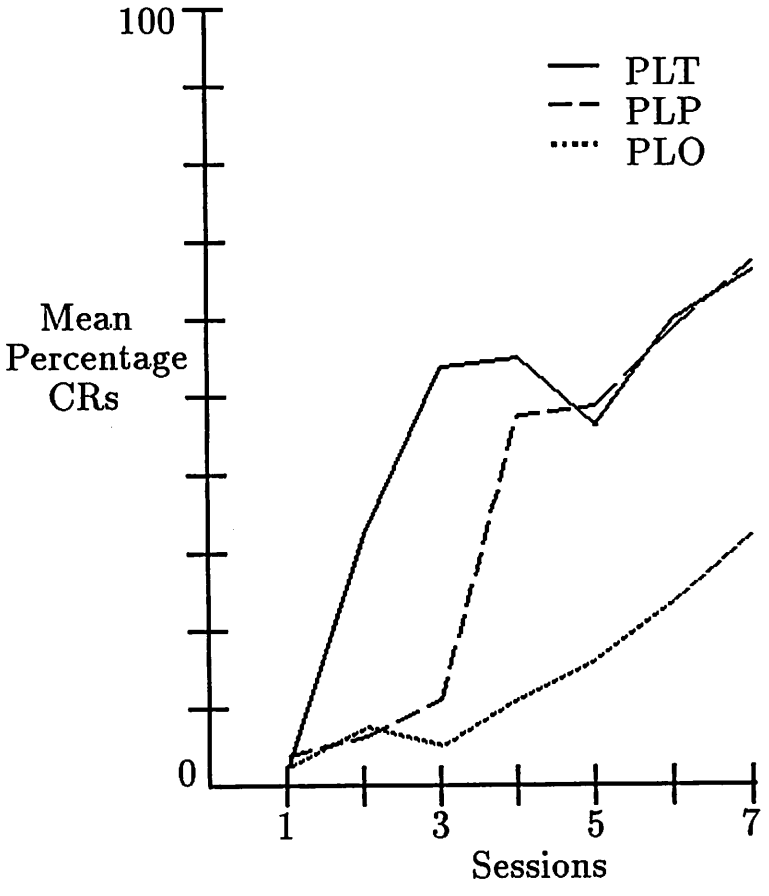

Figure 6.1: **Summary of Experiment 4.** Mean percentage conditioned responding for each group during each day of Experiment 4.

6.1.2 Experiment 4: Results and Discussion

Following acquisition training, mean percent responding over the eight days of training was as follows: group PLT, 55.1%, SE = 6.0; group PLP, 29.4%, SE = 7.4; group PLO, 7.7%; SE = 3.8.

Figure 6.1 shows the percentage conditioned responding obtained for each group during acquisition. The figure suggests that groups PLT and PLP conditioned more rapidly than group PLO. Mean percentage CRs over the seven days were as follows: Group PLT, 45.3%, SE = 4.1; Group PLP, 36.2%, SE = 5.0; Group PLO, 12.8%, SE = 3.3. An analysis of variance was performed on the group mean percentage responding over days. The main effects of Days and Condition were significant

(Days: F(6) = 19.820, $p < .001$; Condition: (F(2) = 7.733, $p < .006$) and the Day x Condition interaction was highly significant (F(2,12) = 3.316, $p < .001$).

Figure 6.1 suggests that Group PLT conditioned more rapidly than did Group PLP. The data were re-analyzed by dividing into two stages: Early Acquisition (defined as days 1 through 3) and Late (defined as days 4 through 6). Each animal thus contributed one score for Early Acquisition and one score for Late Acquisition. For Early Acquisition, Group PLT responded significantly more than did Group PLP (PLT: 29.4%, SE = 6.3; PLP, 6.5%, SE = 4.4; $t(14) = 2.95$, two-tailed $p = .011$). The Groups showed no difference in responding during Late Acquisition (PLT: 52.8%, SE = 9.5; PLP: 54.4%, SE = 8.3; t(14) = .120, two-tailed $p < .91$).

Analysis of trials to criterion yielded further evidence that Group PLT conditioned more rapidly than did other groups. Table 6.1 shows the mean trials required to attain criteria for all groups. Group PLT reached all criteria faster than did the remaining groups. Group PLT was significantly faster than Group PLP on only one criterion, however, that being Trials to First CR (Mann-Whitney U=10, $n_1 = 8$, $n_2 = 8$, two-tailed $p = .02$).

Table 6.1 Mean Trials to Criterion, Experiment 4.

	Criterion		
Group	Trials to First CR	3 CRs in Row	5 CRs in Row
PLT	53.6 ± 11.3	122.1 ± 29.4	128.5 ± 30.3
PLP	147.4 ± 34.2	163.5 ± 34.2	171.9 ± 31.3
PLO	209.4 ± 41.6	270.4 ± 41.6	290.8 ± 32.4

Taken together, Experiments 3 and 4 show that rate of conditioning to a given CS can be facilitated by the addition of another and relatively more salient stimulus during the US. As noted in Chapter 5, the TD and TD_{RTS} models predict this facilitation.

The TD, TD_{RTS}, and SBD models all predict that conditioning is better when CS offset occurs at US onset than when CS and US coterminate. The data of Experiment 4 support this prediction. Extension of the light CS to occur throughout the US (Group PLO) produced dramatically less conditioning than the condition in which light offset coincides with US onset (Group PLO). T-tests of overall mean percentage CRs revealed that Group PLP showed stronger conditioning than did Group PLO ($t(110) = 2.865$; two-tailed $p < .01$).

As mentioned in Chapter 3, a well-documented feature of the rabbit NMR is UR modulation by both neutral and conditioned stimuli (Harvey, Gormezano, and Cool-Hauser, 1985; Hupka et al., 1970; Ison and Leonard, 1971; Weisz and LoTurco, 1988; Weisz and McInerney, 1990; Weisz and Walts, 1990; Young, Cevaske, and Thompson, 1976). Hupka et al., (1970) demonstrated that the amplitude of the UR was greater on paired CS-US trials than during US-alone probes presented during the ITI, even before CRs appeared, and relative to controls that received only paired trials. Their data supported the position that rabbits do not show conditioned diminution of the UR. Ison and Leonard (1971) characterize the course of reflex facilitation and inhibition using a number of parametric manipulations *before* conditioning was manifested. They showed that whether a given CS inhibits or faciliates NMR amplitude depends on stimulus intensity, US intensity, and the temporal relationship between the CS and US. Harvey et al., (1985) and Young et al., (1976) demonstrated a close relation between reflex facilitation and the empirically-observed ISI function. The amount of reflex facilitation elicited by a CS was an accurate indication of the conditionability of that CS (Harvey et al., 1985).

The reflex facilitation data prompted an examination of UR amplitudes in the present experiment. The effects of each of the training procedures on the amplitude of the UR, prior to conditioning, was assessed by examining the first 10 trials of

the first conditioning session for each group.[1] Mean UR amplitudes were as follows: group PLT, 1.0 mm; group PLO, 1.3 mm; group PLP, 2.6 mm. An analysis of variance performed on the mean UR data revealed a highly significant effect of condition (F(2,17), two-tailed $p < .008$). Post-hoc tests of the mean UR amplitudes showed that Groups PLP and PLO were significantly different (t-test, 10 df, two-tailed $p = .008$), but Groups PLO and PLT were not significantly different (two-tailed $p = .578$).

The amplitude of the UR early in conditioning was therefore not a reliable indicator of the results after 7 days of conditioning. Group PLT showed the lowest average UR amplitude early in conditioning, yet proved to be the most readily conditionable. Group PLO showed a slightly higher UR amplitude (though not significantly different from PLT); however, this group conditioned more slowly than the others. Finally, Group PLP displayed the largest UR amplitudes at the start of training, yet conditioned more slowly than group PLT. These results suggest that reflex facilitation did not contribute to the facilitation of responding observed in the compound conditioning group.

6.1.3 Experiment 4: Conclusions

Experiment 4 examined conditioning under the CCO Protocol and relevant single-CS control groups. Group PLT, the CCO Group, received pairings of a light followed by coincident presentation of tone and US. Group PLP received a light that ended before US onset, while Group PLO received light, followed by the US, such that light and US offset coincided. Experiment 4 showed that animals trained under the compound procedure conditioned more rapidly than did single-CS controls. Having replicated the facilitation effect previously observed in Experiment 3, possible mechanisms subserving such facilitation should be considered.

[1] *UR data for one squad of the PLT group was lost due to operator error.*

Reflex Facilitation

Comparisons between the data of Experiment 4 and the UR modulation literature (previous section) are complicated by a number of methodological considerations. For example, the insertion of US and CS alone probes, a typical feature of reflex modulation studies, was not done in Experiment 4. Therefore, there was not baseline against which to measure UR amplitudes when the US was not preceded by the CS. However, parallels based on UR amplitudes from trials early in conditioning suggested that the superiority of conditioning in PLT group was not due to a reflex facilitation. Consistent with the UR data of Experiment 4 is the fact that in Experiment 3, average UR amplitude for group DT, the facilitated group was not significantly greater than that obtained with group D.

Contributions of Associative Processes

It is possible that the pairing of the tone and the US underlies the facilitation of conditioning exhibited by the PLT animals. Consistent with involvement of an association between the tone and the US is the fact that on the first day of training, the performance of all groups was essentially the same. Were non-associative factors to be involved, one might expect that some facilitation would have been observed during the first training session.

Kehoe (1982) reviews a number of studies of serial compound conditioning that are relevant to the present discussion. Kehoe summarizes evidence that the CS2-US interval, the CS1-CS2 interval and CS1-US intervals are important factors in determining the response tendencies for both the compound and for individual elements. Facilitation of conditioning for both the compound and for CS1 probe trials tends to occur when both a) the CS1-US interval is not optimal and b) the CS2-US interval is optimal. In Experiments 3 and 4, facilitation occurred with an optimal CS1-US in-

terval and a non-optimal CS2-US interval. Therefore, the facilitation is not expected based on the existing literature. An important factor in this unusual result may be the fact that CS1 was less salient than CS2. Conditioning is generally not observed in the rabbit NMR with an ISI of 0. It is thus unusual that a CS that cooccurs with a US would facilitate responding to a prior CS. Consistent with this, Kehoe, Feyer, and Moses (1981, Experiment 3A) observed poor conditioning to both CS1 and CS2 in a serial compound having a CS2-US interval of 0, a CS1-US interval of 2,800 ms with CS1 was 400 ms in duration. However, they provided no direct comparison to a single-CS control, did not present data for percentage responding to the compound, and used a CS2 and a US of 50 ms duration. In addition, Kehoe and colleagues counterbalanced the order of stimuli, so that half the group received a tone-light compound while the other half received a light-tone compound. In Experiment 3, as we have noted, the faciliation effect appeared to depend on the salience of the stimulus that overlaps the US. If the stimuli used in the Kehoe, et al., study were not equally salient, then the ability of CS2 to potentiate responding may have been masked.

Further comment should be made regarding the possibility of simultaneous conditioning in the rabbit NMR preparation. Smith et al., (1969) found that a simultaneous conditioning procedure using a 50 ms CS and US resulted in no evidence of conditioning (as assessed by responding during probe trials). However, Burkhardt and Ayres (1978) showed that for conditioned-lick suppression behavior in rat simultaneous conditioning increased as a joint function of CS-US overlap and US duration. This suggests that simultaneous conditioning in the rabbit might occur at longer CS and US durations, like the ones used in the Experiments 3 and 4.

As discussed previously in Section 1.4.1, the controversy continues regarding simultaneous conditioning in a number of conditioning preparations. The facilitation of acquisition of a compound containing a simultaneous CS2-US ISI provides

another, potentially useful, avenue for investigating if and when excitatory conditioning can be expressed in conditioning of the rabbit NMR.

CHAPTER 7

CONCLUDING COMMENTS

7.1 Summary

The present research focused on the abilities of TD, TD_{RTS} and SBD models to account for various features of learning of the conditioned rabbit NMR, and tested certain procedures for which the models predict different behavioral outcomes. A summary of the findings of computational and behavioral investigations of these models is presented below.

7.1.1 Computational Studies

Simulation experiments involving the TD model show that the model scould be modified and extended to encompass aspects of response topography in the classically conditioned rabbit NMR, such as decreasing CR latency and increasing CR amplitude with training. The extended model, denoted as the TD_{RTS} model, includes a modified TD learning rule and incorporates a representation of the CS that is a s template for the desired response topography. The template CS representation had previously been used with success with the SBD model.

Simulation experiments with the TD_{RTS} model assessed parameter sensitivity and its ability to reproduce the set of predictions generated by the TD model. The

TD_{RTS} model proved to be sensitive to the parameters that shape the input trace, such that a rather restricted set of parameter values is required to generate reasonable response topographies and rates of learning given the 10 ms time step assumed for computations. The set of allowable values for parameters of the learning rule is also highly restricted. In particular, certain features of compound conditioning were shown to require the specification of high values for γ. Values outside the optimum of .99 resulted in failure to predict second-order conditioning and certain blocking effects. With the optimal parameter set[1] the TD_{RTS} reproduces the performance of the TD model and generates much of the same topographical features described by the SBD model (Moore et al., 1986), and does so with fewer variables.

The TD_{RTS} model represents an alternative to the SBD model as a description of classical conditioning because ISI functions are generated without postulating a variable decay rate for eligibility or a variable US effectiveness term, mechanisms which are essential for the SBD model's performance. However, the TD_{RTS} model cannot account for the results of Experiment 1, in which trace conditioning with a fixed trace duration was better with a longer than a shorter CS duration (a result that SBD model predicts because of its variable eligibility decay rate), nor can it encompass conditioned diminution of the UR (a result that SBD predicts because of the variable US term, λ'). The TD_{RTS} model could in principle be modified to encompass these results.

7.1.2 Behavioral Studies

A summary of the findings of Experiments 1 through 4 follows.

- Experiment 1 extends the literature on parameters that affect the form of the ISI function for trace conditioning. The data of Experiment 1 suggest

[1]Given the 10 ms time step, the optimal parameter values are as follows: CS representation: $m = .35, h = 1.0, k = .85$; for the learning rule $\gamma = .99, c \leq .1, \beta = .95$.

that CS intensity and duration interact to determine acquisition in a trace conditioning procedure. Experiment 1 provides important support for the variable eligibility decay rate of the SBD model.

- Although only one replication was completed, the data of Experiment 2 show that forward-delay conditioning is not always better than trace conditioning at a given ISI. The trend in those data was for better conditioning with a trace conditioning procedure using an intense CS: the offset of an intense stimulus, when optimally located with respect to the US, is a salient CS in its own right.

- Experiment 3 was designed to assess the CCO procedure, for which the TD and SBD models provide differing predictions. In the CCO procedure, two CSs, CS1 and CS2, are presented. CS1 is followed by a simultaneous compound of CS2 and US. Results of the CCO groups are compared against CS1 controls. Experiment 3 showed that when CS1 was a relatively less salient stimulus than CS2, conditioning was facilitated for the compound relative to the CS1 control. This result provided support for the TD model. Weak evidence that CS2 was a conditioned inhibitor suggested that some support for the SBD model as well.

- Experiment 4 replicated the facilitation effect. The results suggested that the facilitation was due to the simultaneous presentation of a strong CS and the US. The possibility exists that transfer of excitatory conditioning occurred from the simultaneous CS to the first CS in the compound.

The behavioral studies of this dissertation suggest that further empirical investigations should be conducted to assess the descriptive power of the SBD, TD, and TD_{RTS} models.

7.2 Future Directions

7.2.1 Investigations of CS Representation

Both the behavioral work and the computational work point out the need for further investigation of the nature of CS representation. The results of Experiments 1 and 2 call for a further examination of the effects of CS parameters on the ISI function. Experiments 3 and 4 showed a facilitation of conditioning that occurred when a relatively *salient* CS overlapped the US.

The SBD and TD_{RTS} models do not address parameters of the CS such as stimulus intensity or salience. The findings of the simulation experiments reported here for the TD_{RTS} model and elsewhere for the SBD model (Blazis and Moore, 1987) indicate that the template representation in its current implementation cannot be changed without affecting the models' predictions for other conditioning protocols. However, all of the models considered here could in principle be extended to phenomena as stimulus intensity and salience effects. Therefore, alternative representations must be considered.

The strategy for developing more useful CS representations involves consideration of behavioral data, physiological data, and CS representations that have been proven to be useful in other models. Behavioral data like that obtained in experiments 1 and 2 suggest that stimuli other than the nominal CS are used by the animal to generate CRs. For example, the prevalence of CR-onset latencies during the offset of the trace CS in Experiment 2 suggests that the offset of the CS delineates the trace interval itself as a CS.

A growing literature pertaining to physiological investigations of neuronal activity in sensory and limbic regions during conditioning could provide important clues about the nature of adaptive stimulus representations in classical conditioning (for example, Berger, Weikart, Bassett, and Orr, 1986; Gabriel, Miller, and Saltwick,

1976; Romano, Steinmetz, Mikhail, and Patterson, 1986; Ryugo and Weinberger, 1978). These studies indicate changes in neuronal activity, correlated with learning, that might have implications for the development of a normal CR.

Finally, CS representations that have proven useful for other models should be investigated with respect to the SBD and TD_{RTS} models. One reasonable avenue for future computational experiments of the TD model would be to investigate the model's performance when coupled with a distributed CS representation. To overcome certain limitations of the template representation used in prior models, the Desmond-Moore model (Desmond, Blazis, Berthier, and Moore, 1986; Desmond and Moore, 1988; Desmond, 1988) assumed that CS input sets up a stimulus trace that has a timing function which is a property of the architecture of sensory elements. In this representation, the input of the CS on a given time step consists of the pattern of activation over elements arranged in a tapped delay-line or transversal filter. Moreover, separate tapped delay-lines are initiated by CS onset and CS offset. Preliminary work by Sutton and Barto (1990, in press) suggests that such a CS representation could allow the TD model to generate weights that increase to peak prior to the US, a result that contrasts with the decreasing intratrial weights that can result with the template representation.

7.2.2 Simultaneous Conditioning

The behavioral and computational work also point to further investigation of the nature of a simultaneously conditioned CS in rabbit NMR conditioning. The TD_{RTS} and SBD models disagree completely as to whether a simultaneously conditioned CS is a conditioned excitator or a conditioned inhibitor, yet neither model fully accounts for the facilitation effect.

Future investigations suggested by the data of Experiments 3 and 4 include second-order conditioning studies that examine the possibility that a simultaneously-

conditioned CS can serve as a secondary reinforcer in the conditioned NMR. Rescorla (1980) suggested that the extent to which a CS can serve as a secondary reinforcer is an indication of the associative character of that CS. Also, a modified version of Experiment 3 could shed further light on the possibility that the overlapping CS can have dual characteristics; an excitatory stimulus-stimulus association but an inhibitory stimulus-response relation. A modified version of that experiment might use CSs from three modalities (visual, auditory, and tactile) instead of the two that were used, in order to reduce stimulus generalization.

7.2.3 Implementation of the TD Model in Real Neural Networks

This dissertation shows that further development of the TD and SBD models is dependent on explorations of useful CS representations and the acquisition of further behavioral data. This research does not evaluate the SBD and TD models as descriptions of how cells or small networks learn. The SB and SBD models have been established as fertile ground for theoretical explanations of how cells or small networks might be classically conditioned (Barto and Sutton, 1982; Moore and Blazis, 1989; Sutton and Barto, 1981). As dictated by the strategy shown in Chapter 1, future research should address the possibility of extending the TD model in this direction. Features of the TD model that are theoretically interesting for the biological network researcher include the following:

- The TD model specifies that learning involves an anticipation of activity along a particular pathway, that being the input pathway of the US. This contrasts with the SB and SBD models, which view learning as adjusted by activity along output pathways. A neural implementation of the TD model and variants like the TD_{RTS} model may require a specific US input pathway.

- The discrepancy used in the learning rule involves $P(V(t), x(t))$ and $P(V(t), x(t-1))$, which means that the unit or network must retain the input of the previous time step in some way in order to combine it with $V(t)$. In the SB and SBD models, the type of summation represented by P is always done with time indices that match; these summations have been theoretically aligned with network architectures (Barto and Sutton, 1982; Moore and Blazis, 1989). Barto (1990, personal communication) has indicated that with a sufficiently small learning rate, the Ps in TD model can be computed with matching time indices. In any case, the difference in the way these error terms are computed by the models may suggest different neural network architectures.

- Once the size of the time step is established, the TD model is very sensitive to the value of parameter γ, both in its predictions for topography and for overall results of conditioning protocols (i.e., Chapter 5). In previous work, Moore and Blazis (1989) exploited parameter sensitivity as a means of specifying networks in cerebellum that could compute the SB learning rule. Any network that computes the TD model should account for or be built upon this aspect of the TD model's parameter sensitivity.

7.3 Model Evaluation: Other Real-Time Models of Conditioning

The TD_{RTS} model fares at least as well as the SBD model has in describing many aspects of classical conditioning, and does so without the addition of the additional parameters represented by λ' and the decay of $\bar{x}$ currently required by the SBD model. Both models are compromised by the representation of the CS. This section briefly considers the validity of future studies of the SBD and TD models when considered in light of other real-time models of the classically conditioned NMR.

Both the SBD model and the Desmond-Moore model mentioned previously address the problem of response topography. Unlike the SBD model, in the Desmond-Moore model onset of a CS initiates two trace processes, an onset and an offset process. Two adaptive elements modify the weight of each of the input elements with respect to the US. One adaptive element computes an expectation of the temporal locus of the US, while the second adaptive element computes associative strength between CS and US, based in part upon the value of US expectation. The Desmond-Moore model is a network model: Its ability to correctly predict response topography for trace and long-CS duration delay conditioning procedures basically arises from its network architecture. The model also generates complex waveforms like those observed by Millenson, Kehoe, and Gormezano (1977). However, the model requires additional units in order to describe higher-order conditioning, a feature of classical conditioning yielded by the algorithms of the SBD and TD models without a need for additional assumptions or algorithms.

The SB, SBD, and TD models do not provide an account of the totality of classical conditioning phenomena, for example attentional phenomena such as latent inhibition.[2] Schmajuk and Moore (1988) propose a model (Schmajuk-Pearce-Hall, or SPH) that is capable of describing many features of classical conditioning, including higher-order conditioning, sensory-preprocessing, and latent inhibition. The model has been applied, with some success, to the description of effects of hippocampal lesions. The SPH model also describes NMR topography by imposing performance rules upon model output, rather than presenting the adaptive element with a template of the to-be-learned response. However, these performance rules do present problems for the model. Namely, the rules result in inappropriate timing of the CR peak for ISIs of less than 200 ms duration (Schmajuk, 1986; Schmajuk and Moore, 1988).

[2]Sutton and Barto (1987) have noted that several connectionist learning mechanisms may be combined with the TD model in order to extend its range of classical conditioning predictions.

The Desmond-Moore and SPH models are attractive alternatives to the SBD model with regard to describing classical conditioning of the rabbit NMR. The SPH model in particular encompasses an impressively large array of the data. However, it is also very complex, requiring the application of most of twenty equations per time step. The Desmond and Moore model requires a relatively large number of units. By improving CS representations, it is possible that either the SBD model or the TD model could expand their predictive repertoires and remain fairly parsimonious descriptions of classical conditioning.

APPENDIX A

THE SUTTON-BARTO-DESMOND MODEL

Moore et al., (1986) extended the SB model to account for CR topography by treating the input of the i^{th} CS in the model, x_i, as a function of time. In the SBD model, each time step t corresponds to 10 ms. At CS_i onset, $t = 0$. For time steps $t = 1, ..., 7$, the input $x_i = 0$. For $t > 7$ and until CS_i offset, x_i is defined as follows:

$$x_i(t) = [tan^{-1}(mt - 5.5) + 90]/180\,h. \tag{A.1}$$

The parameter m, $m > 0$, controls the rate of rise of x_i, and the parameter h, $h > 0$, controls CR amplitude. Blazis and Moore (1987) showed that the optimal values of these parameters, that is, the values that produce the best joint predictions for response topography and conditioning, are $m = 0.35$ and $h = 0.75$ or 1.00.

A second function returns the CR generated by x_i to its pretrial baseline. It is implemented at CS_i offset, decays geometrically, and is computed as follows:

$$x_i(t) = kx_i(t-1), \tag{A.2}$$

where $0 < k < 1$. In the SBD model, the optimal value for k is 0.85.

ΔV_i is computed as in the SB model:

$$\Delta V_i(t) = c[s(t) - \bar{s}(t)]\bar{x}_i(t), \tag{A.3}$$

where c is a learning rate parameter, $0 < c \leq 1$, $s(t)$ is the element's output at time step t, and $\bar{s}(t)$, defined below, is a function of $s(t)$ from preceding time steps. According to Moore et al (1986), the optimal value of $c = 0.15$.

The output of the learning element at time t, denoted $s(t)$, is defined as:

$$s(t) = \sum_{i=1}^{n} V_i(t)x_i(t) + \lambda'(t). \tag{A.4}$$

In the context of the SBD model, $s(t)$ is a measurement of the displacement of the NM during the course of the trial. While s can take on any real value in the SB model, in the SBD model it is confined to the closed unit interval. Limitations on s are imposed because of physiological constraints of the NMR.[1]

Variable $\lambda'(t)$ in Equation A.4 equals 0 prior to the occurrence of the US. During US presentation λ' is calculated as the difference between λ, the weight of the US when all $V_i \leq 0$, and the largest positive starting weight among all CSs present on a given trial. (Starting weight refers to $V_i(0)$). Thus, if V_i is the largest starting weight among the CSs present on the trial, while the US is present[2]

$$\lambda' = \begin{cases} \lambda - V_i & \text{if } 0 \leq V_i \leq \lambda; \\ 0 & \text{if } V_i > \lambda; \\ \lambda & \text{if } V_i \leq 0. \end{cases} \tag{A.5}$$

At US offset, λ' decreases as follows:

$$\lambda'(t+1) = 0.9\lambda'(t). \tag{A.6}$$

The prediction of output for the current time step, denoted $\bar{s}$, is computed by:

$$\bar{s}(t) = \beta\bar{s}(t-1) + (1-\beta)s(t-1), \tag{A.7}$$

where $0 \leq \beta < 1$. The parameter β determines the rate of decay of $\bar{s}$. The optimal value for β is 0.6.

Variable $\bar{x}_i(t)$ in Equation A.3 is a CS_i duration-dependent stimulus trace that defines the period and extent to which the i^{th} connection is eligible for modification. For a given time step t, the eligibility trace is defined as follows:

$$\bar{x}_i(t) = x_i(t-3), \tag{A.8}$$

[1] Negative values of s imply nictitating membrane retraction and exopthalmus, which are not typically observed in the rabbit. The upper bound of 1.0 reflects the fact that although the amplitude of the NMR is directly related to the intensity of the eliciting stimulus, there are limits on the number of involved motoneurons and their rate of firing (Berthier and Moore, 1980; Moore and Desmond, 1982).

[2] Variable λ' implements the idea that US effectiveness can diminish progressively with training (e.g., Donegan and Wagner, 1987; Mackintosh, 1983). Were $\lambda'(t)$ to be replaced in Equation A.4 by λ, post-US computations would cancel increments in V_i during time steps preceding US offset, and as a result there would be no net learning (see Moore et al., 1986). In addition, response topography would be compromised.

during time steps t that the CS_i is on.[3] $\bar{x}_i$ begins its decline three time steps after CS_i offset:

$$\bar{x}_i(t) = \delta \bar{x}_i(t-1), \tag{A.9}$$

where $\delta = e^{-3/d}$, $d = \max\{25, CS_i \text{ duration in units of 10-ms}\}$. The computations shown define a period of eligibility which begins some time after CS_i onset and persists beyond CS_i offset.[4]

The eligibility trace, $\bar{x}$, and in particular its decay, is responsible for the ability of the SBD model to generate ISI functions that are consistent with the behavioral literature. For the rabbit NMR, the ISI function is concave-down with a peak at an optimal ISI of 250 to 350 ms (Schneiderman and Gormezano, 1964; Schneiderman, 1966)[5]. With a relatively limited number of trials the SBD model predicts a concave-downward ISI function. However, like the TD and TD_{RTS} model, the SBD model predicts that conditioning is roughly decreasing function of ISI (when CS Duration = ISI) at asymptote (after roughly 350 trials). In Table A.1, an asymptotic ISI function with ISI = CS duration is shown ($\lambda = .9$, US duration = 50 ms).

Table A.1 Predicted Asymptotic ISI Function with the SBD Model

ISI in ms	100	200	300	400	500	600	700	800	900	1000
V_i	.825	.702	.672	.603	.547	.502	.464	.431	.403	.379

[3]In earlier descriptions of the model, for example Moore et al., (1986), the lag was stated to be 3 time steps, or 30 ms, but the code generating the SBD simulator actually produced a lag of 4 time steps or 40 ms.

[4]In a *forward-delay procedure*, the onset of the CS precedes that of the US; the US typically occurs either during the CS or immediately after CS offset. In a *trace conditioning procedure*, the offset of the CS occurs some time before the onset of the US. The period of time intervening between CS offset and US onset is called the trace interval.

[5]Without the duration-dependent decay mechanism shown in Equation A.9, the SBD model predicts that V increases with ISI.

APPENDIX B

SIMULATION SOFTWARE

The computational experiments in this dissertation were run with GEN_SIM, "General Simulator" software written in C by Diana Blazis to run under the Sun-OS operating system. GEN_SIM computes input, output and update rules for models of adaptive behavior. The program can be run interactively or automatically ("batch" processing). In interactive mode the program provides popup menus which allow the user to change learning algorithms, model parameters, and learning procedures at will (Figure B.1). The program displays the appearance of the paradigm, computes the experiment, and plots model output in several forms (Figure B.2). In batch mode, the program reads prepared command files and automatically generates hardcopy graphics output on request.

Owing to the size of the simulator, the code comprising the simulator is not provided. Queries about GEN_SIM should be directed to the author.

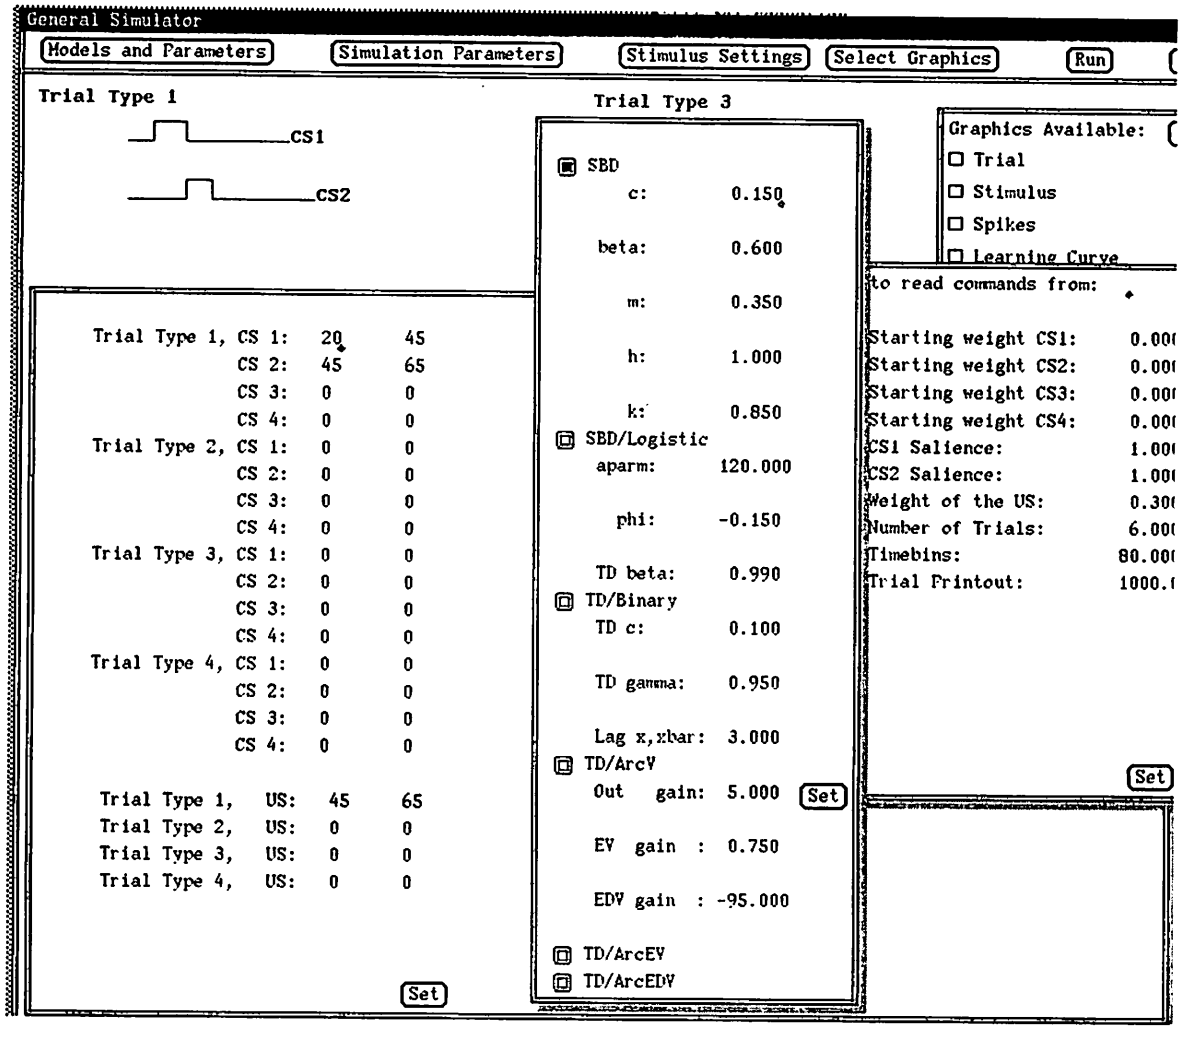

Figure B.1: **GEN_SIM: User Interface.** Depicted above are activated pop-up menus from the interface of GEN_SIM. These popups allow selection of stimulus timing, model parameters, simulation parameters, and graphics output.

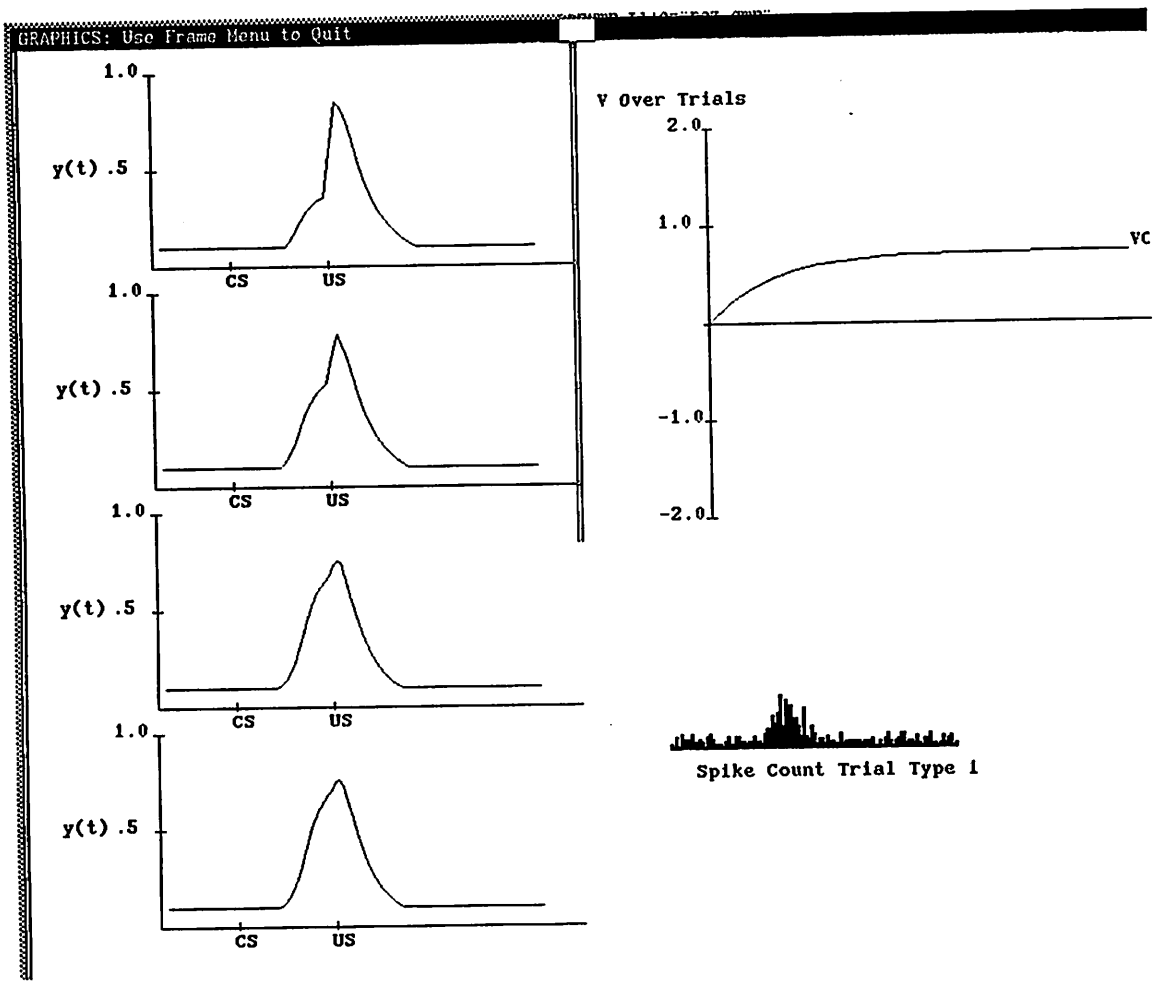

Figure B.2: **GEN_SIM: Sample Simulator Output.** Depicted above are samples of representations we use to evaluate outputs of computational models. Left: Representation of a typical CR for a given CS over trials. Upper Right: Learning Curve plotted as V over trials. Lower Right: Simulated neuronal firing.

APPENDIX C

DERIVATION OF TD_{RT} RULE

For a fuller explication of the material for this appendix, consult Sutton and Barto (1990, in press).

The TD model assumes that an organism in a classical conditioning situation tries to predict future values of the US signal, λ. Let $\bar{V}(t)$[1] denote the prediction at time t of all future values of λ:

$$\bar{V}(t) = \lambda_{t+1} + \lambda_{t+2} + \lambda_{t+3} + \lambda_{t+4} + \cdots \tag{C.1}$$

The assumption that later values of λ are discounted is implemented by assuming that each of the future values of λ is multiplied by some fraction γ, $0 < \gamma < 1$, for each step that λ is delayed. Sutton and Barto assume the weighting occurs according to a power function so that one-step delayed λ is discounted by γ, two-step by γ^2, etc. The weighting could occur by some other function, but the exponential was selected due to computational convenience.

$$\bar{V}(t) = \lambda_{t+1} + \gamma\lambda_{t+2} + \gamma^2\lambda_{t+3} + \gamma^3\lambda_{t+4} + \cdots \tag{C.2}$$

Expressed as the exponential function in Equation C.2, $\bar{V}(t)$ constitutes a prediction about the future area, or alternatively a weighted sum, of future US input. For the purposes of the present study, a prediction about the average amplitude of the US was desired. To convert from a weighted sum to a weighted average, the coefficients (γs) must sum to 1. Since the limit of the series shown in Equation C.2 is (1/1-γ), the appropriate multiplier for each coefficient is (1-γ). Rewriting

[1] $\bar{V}$ is equivalent to P(v,x) defined in Chapter 1.

Equation C.2 to include the coefficient yields the following:

$$\bar{V}(t) = (1-\gamma)\lambda_{t+1} + (1-\gamma)\gamma\lambda_{t+2} + (1-\gamma)\gamma^2\lambda_{t+3} + (1-\gamma)\gamma^3\lambda_{t+4} + \ldots \quad \text{(C.3)}$$

or alternatively,

$$\bar{V}(t) = (1-\gamma)\lambda_{t+1} + \gamma(1-\gamma)[\lambda_{t+2} + \gamma\lambda_{t+3} + \gamma^2\lambda_{t+4}\ldots] \quad \text{(C.4)}$$

The quantity in the brackets is $\bar{V}(t+1)$, that is, if Equation C.3 were rewritten for $\bar{V}(t+1)$ the result would be

$$\bar{V}(t+1) = (1-\gamma)\lambda_{t+2} + (1-\gamma)\gamma\lambda_{t+3} + (1-\gamma)\gamma^2\lambda_{t+4} + \ldots \quad \text{(C.5)}$$

or alternatively

$$\bar{V}(t+1) = (1-\gamma)(\lambda_{t+2} + \gamma(\lambda_{t+3}) + \gamma^2\lambda_{t+4}\ldots) \quad \text{(C.6)}$$

Substituting $\bar{V}(t+1)$ into Equation C.4 yields

$$\bar{V}(t) = (1-\gamma)\lambda_{t+1} + \gamma\bar{V}(t+1). \quad \text{(C.7)}$$

The discrepancy term of the model is thus

$$(1-\gamma)\lambda_{t+1} + \gamma\bar{V}(t+1) - \bar{V}(t). \quad \text{(C.8)}$$

APPENDIX D

SIMULATIONS: TEMPLATE REPRESENTATION

D.1 Simulations Set 1A

The following simulation experiments show the effects of varying parameters of the template representation x of the TD_{RTS} model. These parameters include m, which determines the rise of x; h, which determines the amplitude of x; and k, which determines the rate at which x declines after the CS. The values selected for simulation were shown by Blazis and Moore (1987) to noticeably affect the SBD model's behavior. Thus four values of h (0.25, 0.75, 1.0, 2.0), three values of k (0.25, 0.65, 0.85), and three values of m (0.1, 0.35, 0.7) were examined separately.

Presented in this appendix are the results of simulations of acquisition, ISI functions, and response topography.

D.1.1 Learning Curves: Shape and Rate of Acquisition

Learning curves depicted in Figures D.1 and D.2 plot the acquisition of V_i over 300 training trials under variation of the x-shaping parameters (see Figure captions for timing details). For the rabbit NMR, empirically observed learning curves tend to be S-shaped. Like the SBD model and other models of learning, the learning curves generated by the TD_{RTS} model tend to be negatively accelerated functions of trials.

Figure D.1 shows the only parameter that significantly alters rate of conditioning. With $m = 0.1$ rate of conditioning is very slow. For the 300 ms CS illustrated in Figure D.1, stable weights are not attained for thousands of trials. Though not

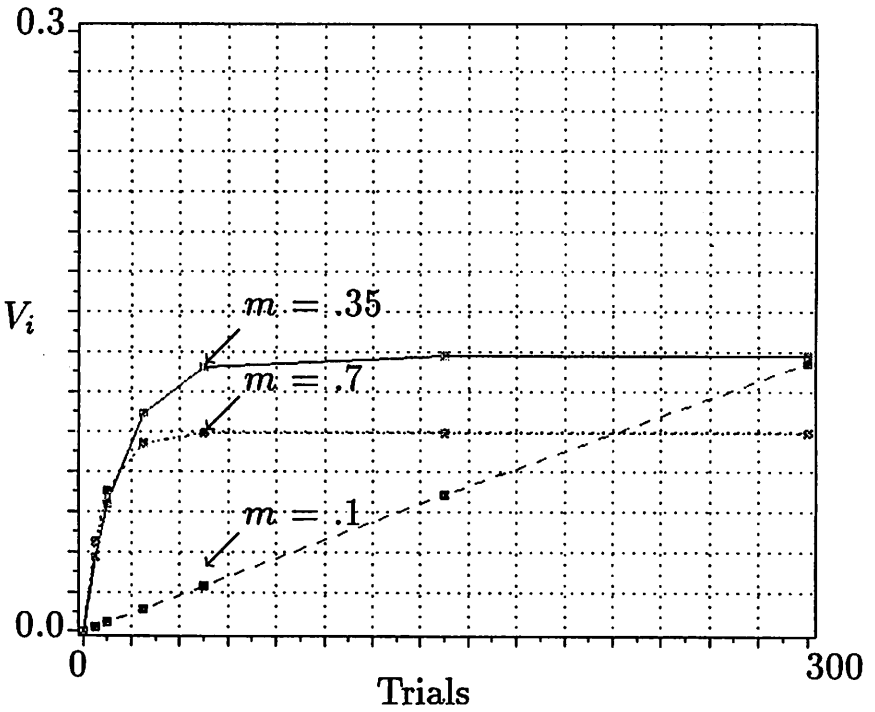

Figure D.1: **Effects of Parameter m on Acquisition.** Data shown were obtained using a 300 ms CS followed by a 50 ms US. $\lambda = .7$.

shown, with $m = .1$ asymptotic level of V_i is well over the expected value predicted by Equation 2.3; for the 50 ms US of intensity .7, the expected maximum value of V_i is about .3. This example shows a kind of input for which the model's actual prediction substantially overshoots the theoretical prediction.

Figure D.2 shows the effects of variation in parameter h on rate of acquisition. By Equation A.1, parameter h controls the amplitude of x, such that amplitude of x increases as h decreases. A consequence of increasing x amplitude is a *decrease* in asymptotic levels of V_i.[1] The decrease in V_i seen with larger amplitudes of x occurs because of the following features of the TD model. In the TD model the static CS generates unweighting during its presence, depending on the value of γ (a point discussed more fully in Chapter 2). Recall that eligibility for changes in V_i depends upon the value of the input x. When signal amplitude is relatively high and stable, eligibility for unweighting is also high, leading to substantial decrements in V_i during the CS.

[1]This result stands in direct contrast to that obtained with the SBD model (Blazis and Moore, (1987), Experiment 1). Increasing x amplitude by decreasing h leads to more rapid rates of acquisition and higher asymptotic values for V_i.

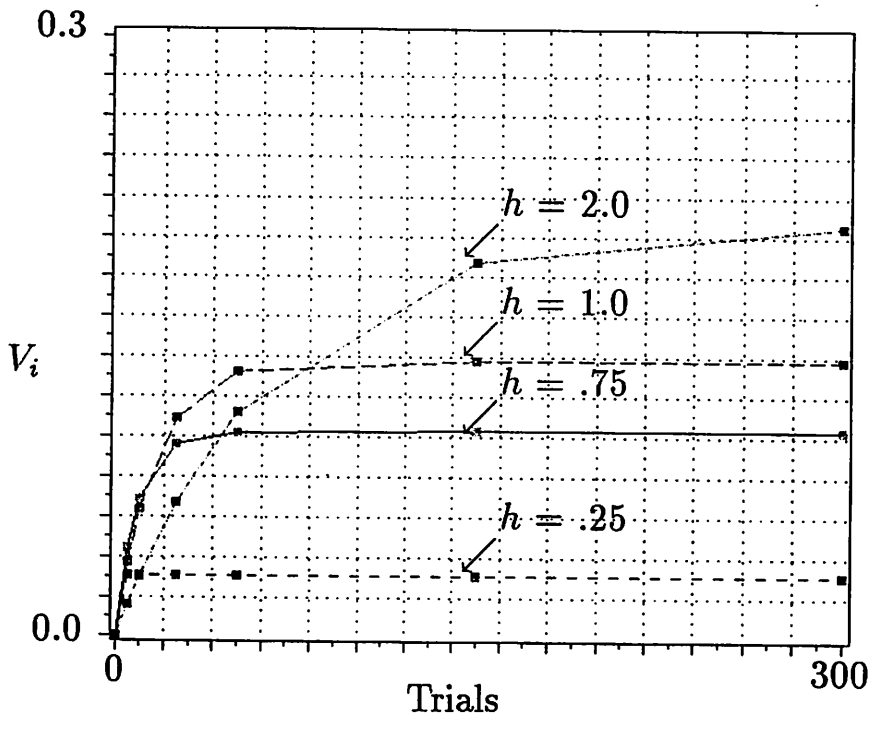

Figure D.2: **Effects of Parameter h on Acquisition.** Data shown were obtained using a 200 ms CS followed by a 50 ms US. $\lambda = .7$.

Learning curves are relatively insensitive to the value of parameter k (not shown). As k increases, asymptotic values of V_i decrease very slightly.

D.1.2 ISI Effects

As noted in Chapter 3, reasonable predicted ISI functions are those that are concave-downward in shape after a relatively small number of trials. Given that the TD and TD_{RTS} models will tend to predict that conditioning decreases with ISI *in the limit*, these studies emphasized selection of x-shaping parameter values that yield reasonable acquisition rates and topography. However, parameter values that aversely affect these features of conditioning can yield some unusual results for ISI functions also.The following experiments consider the form of the ISI function after a small number of trials.

Figures D.3 shows ISI functions obtained after 100 trials with various values of m. When $m = 0.1$, the resulting function is unrealistic because V_i increases uniformly with ISI. This occurs because a) low m slows rate of acquisition and b) acquisition in the model is generally more rapid for long rather than short ISIs. After asymptotic

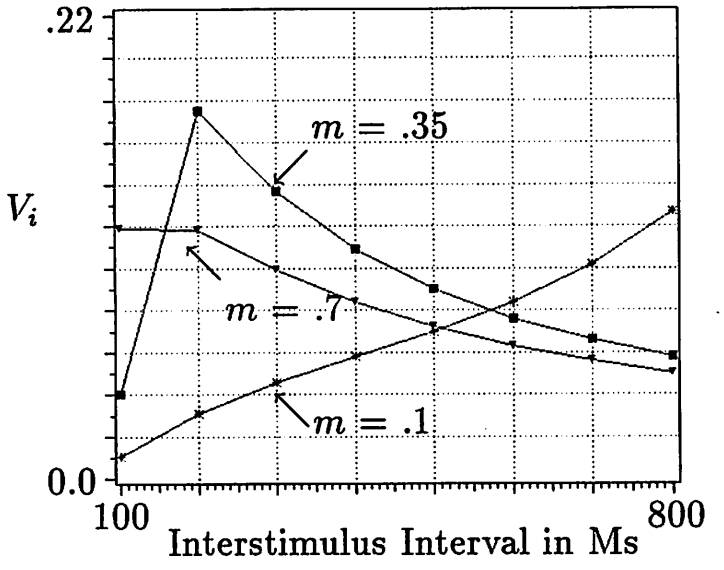

Figure D.3: V_i **as a Function of ISI and** m **in the Forward-delay Paradigm.** Data shown were obtained after 100 trials using a 50 ms US, and with m ranging from 0.1 to .7. $\lambda = .7$.

training, V_i is a decreasing function of ISI since long ISIs reach asymptote at lower levels of V_i.

Figure D.4 shows ISI functions obtained under various values of h. After 100 trials, V_i obtained for $h = 0.25$ is a decreasing function of ISI. When $h = 2.0$, the peak of the ISI function is shifted over to 300 ms rather than 200 ms; this reflects the fact that acquisition is not yet complete for ISIs of less than 300 ms. Were training to continue V_i would be a decreasing function of ISI for all ISIs considered.

Figure D.5 shows ISI functions obtained under variations in parameter k. Its effects on the ISI function are negligible, as indicated by the similarity of the plots.

A point that should be emphasized is that when ISI = CS duration, and when CS-US pairings are given until learning is asymptotic, the TD_{RTS} model predicts that conditioning is a decreasing rather than a concave-downward function of CS duration for ISIs greater than or equal to 100 ms (the shortest ISI simulated here).

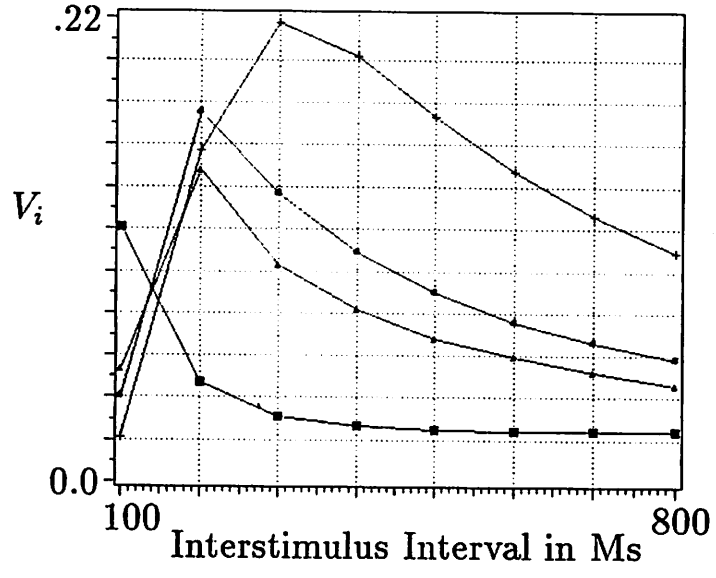

Figure D.4: V_i **as a Function of ISI and** h **in the Forward-delay Paradigm.** Data shown were obtained after 100 trials using a 50 ms US, and with h ranging from .25 (lowest plot) to 2 (highest plot). $\lambda = .7$.

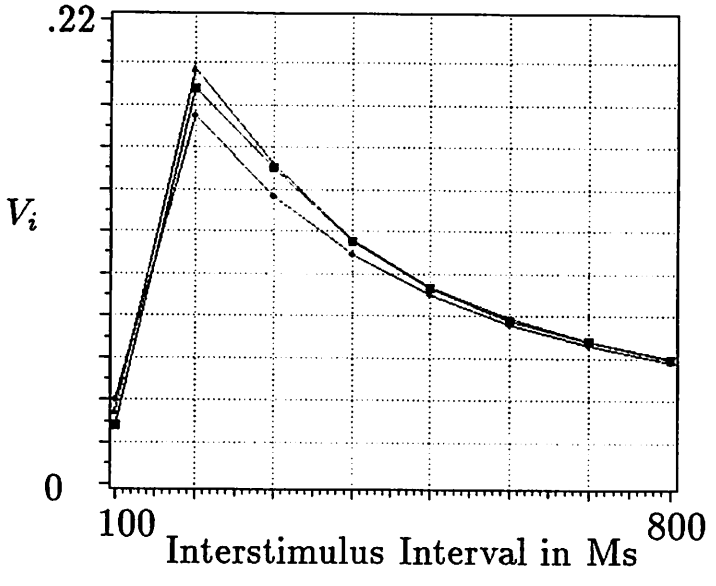

Figure D.5: V_i **as a Function of ISI and** k **in the Forward-delay Paradigm.** Data shown were obtained after 100 trials using a 50 ms US, and with k ranging from .25 (lowest plot) to .85 (middle plot). $\lambda = .7$.

D.1.3 Response Topography

Though ISI functions in the preceding section were generated using 100 trials, response shapes in this section were generated following asymptotic training. This strategy was used in order to make the following points: even at asymptote, low values of m are not suitable for generating response topography, but all values of h generate roughly the same response topography.

Figure D.6 shows simulated responses obtained after asymptotic training for different values of m and for two ISIs. Of the values examined, only $m = .35$ allowed the TD model to produce a realistic response at the 300 ms ISI (Panel A). Since m influences the rate at which x increases, for 100 trials low values of m result in relatively low-amplitude CRs (not shown) while high values result in CRs that rise rapidly and abruptly change slope during the ISI. At asymptote, however, $m = 0.1$ and 0.7 fail to produce CRs that have the smoothly-increasing, S-shaped patterns typical of the conditioned NMR. With $m = .1$, the high value of V (2.2 at trial onset) is responsible for the sharp upturn in the CR. For $m = 0.7$, asymptotic V is very low, but x increases sharply; by Equation 2.4 this results in a square-shaped CR.

Low values of m yield a CR profile that increases over the ISI when longer CS durations are used. Figure D.6, Panel B shows response topographies obtained with a 650 ms CS. Values of $m > .1$ result in CR profiles that decline prior to the US. With $m = 0.1$, however, a CR that increases over the ISI is obtained. Initial portions of the CR show an abrupt change of slope.

Figure D.7 shows simulated responses obtained after asymptotic training for different values of h and for two ISIs. With the short ISI of 300 ms, variations in h have almost no effect on the CR. This is because parameter h has a linear effect on V in the model. Though x on every time step is lower for $h = 2$ than for $h = 0.25$, the asymptotic weight is proportionally higher. The net effect is that CR topography is virtally unchanged by the size of h. However, for a low, fixed number of trials differences would be evident because rate of learning is also affected by h (Figure D.2).

For the 650 ms CS (Figure D.7, Panel B), however, the response profile produced with $h = 0.25$ seems quite different from those produced with the other values of h. With $h = 0.25$, x peaks and levels off early in the ISI. Whenever inputs cease to increase during the ISI, V declines. The remaining values of h produce an input that is increasing, albeit only moderately, during latter portions of the ISI. The profiles shown in Panel B do begin to decline before the US, but the magnitude of this decline is less with $h > .25$.

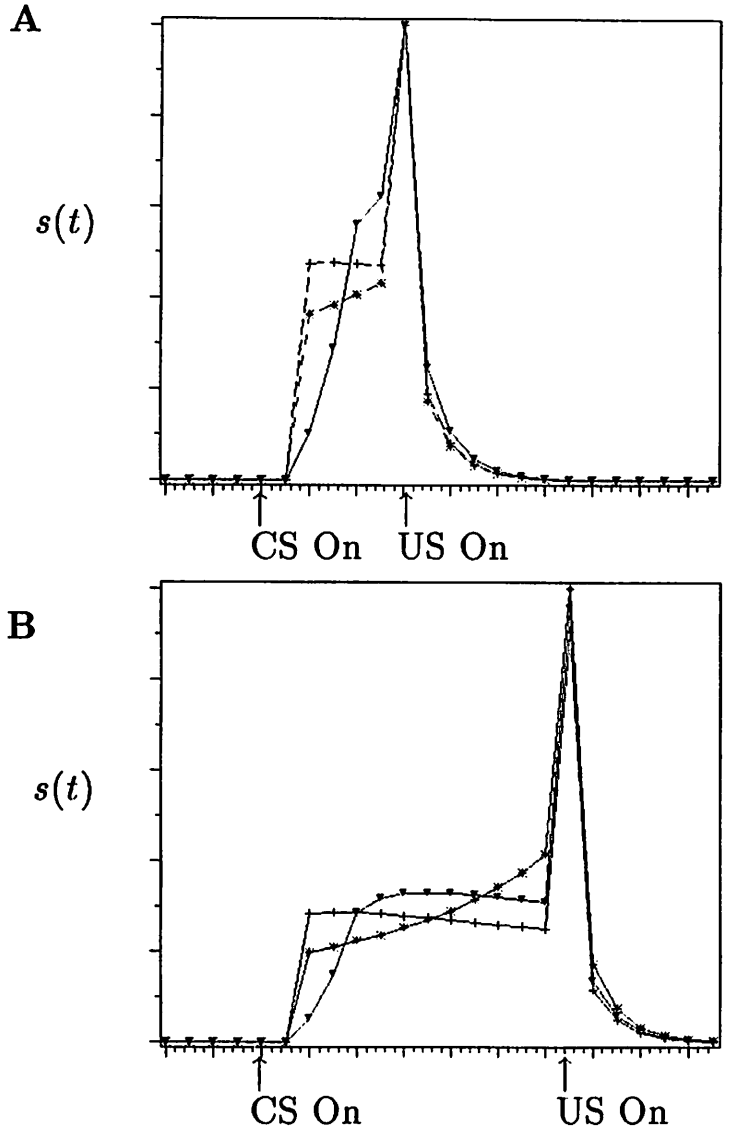

Figure D.6: **Response Topography as a Function of m.** Data shown were obtained at asymptote using a 50 ms US and ISI equal to CS duration. $\lambda = .7$. The starred plot is obtained with $m = .1$, triangles demarcate the plot obtained with $m = .35$, and crosses indicate $m = .7$. In Panel A, the CS duration is 300 ms. Values of V are as follows: for $m = 0.1$, $V = 2.2$; for $m = .35$, $V = .134$, and for $m = .7$, $V = .099$. In Panel B, the CS duration is 650 ms. Values of V are as follows: for $m = 0.1$, $V = 2.2$; for $m = .35$, $V = .134$, and for $m = .7$, $V = .099$.

D.2 Simulations Set 1B

Simulations Set 1B investigated combinations of x-shaping parameter values, m, h, and k. In these simulations, four values of h (0.25, 0.75, 1.0, 2.0), three values of k (0.25, 0.65, 0.85), and three values of m (0.1, 0.35, 0.7) were covaried orthogonally. Experiments assessed the effects of covariations of parameters on terminal levels of V_i, rates of learning and response topography for ISIs of 300 and 650 ms.

D.2.1 Learning Curves: Shape and Rate of Acquisition

Figure D.8 plots V_i as a function of the values of parameters m and h; k is fixed at 0.85. As expected from Simulations Set 1, V_i after 300 trials tended to increase with h and decrease with m. Combinations that included $m = 0.1$ and extreme values of h were exceptions. When $m = 0.1$ and $h = 0.25$, V_i tends to "blow up", that is, the value of V_i is very large. When $m = 0.1$ and $h = 2.0$, terminal levels of V_i are quite large as well. In Figure D.8, this fact is not evident because rate of acquisition for this parameter value combination is slow; at 300 trials the resulting V_i is only about 10 percent of its terminal value of 2.0. Although not shown, the surface shown in Figure D.8 is also obtained when $k = 0.25$ and $k = .65$.

Figures D.9 and D.10 plot V_i as a function of trials for combinations of m and h; k is fixed at 0.85. Figure D.9 was generated with an ISI of 300 ms; Figure D.10 was generated with a 650 ms ISI. The learning curves obtained are generally negatively-accelerated functions of trials. Rate of acquisition is lowest with $m = 0.1$. Most of the data plotted in these figures are consistent with the results of Simulation Set 1; a notable difference is that the tendency of $m = 0.1$ to lead to blow-ups of V_i is curbed when $h = 0.25$, resulting in relatively high weights that could be of use in alleviating the model's tendency for low amplitude CRs.

D.2.2 ISI Effects

Figure D.11 shows value of V_i after 300 trials for a 300 ms (Panel A) and a 650 ms CS (Panel B) with variations in m and h (k is fixed at .85). With $m \geq .35$, V_i is higher for the 300 ms CS than for the 650 ms CS for all values of h, consistent with the literature. With $m = 0.1$, the ISI function is realistic only with $h = 0.25$. With $m = 0.1$ and $h > .25$, the 650 ms CS duration yields a higher V_i than does the 300 ms CS after 300 trials.

D.2.3 Response Topography

Figures D.12 and D.13 depict response profiles under combinations of x-shaping parameter values with ISIs of 300 and 650 ms, respectively. Since negative reinforcement grows as CS presence continues, response characteristics generated by a given parameter combination with the 300 ms ISI (Figure D.12) are not typically maintained for the longer ISI (Figure D.13).

For the 300 ms CS, only combinations involving $m = .35$ resulted in realistic response topographies, and of these, most realistic profiles were obtained with $h < 2.0$. When $m = .1$, response amplitude increases during the CS, but does so slowly and in a linear rather than s-shaped fashion. When $m = .7$, the response rises rapidly but then decreases over time steps leading up to the US, especially with $h = .25$.

For the 650 ms CS, combinations involving $m = .35$ and $h > .25$ resulted in response topographies that rose smoothly, but which tended to decrease slightly in time step prior to the US. The profile generated with $m = .35$ and $h = 0.25$ shows a pronounced decline prior to the US. With $m = 0.1$, a profile that increases over the ISI zoccurs but initial portions of the response are marked by an abrupt change in slope that is not typically observed. Combinations involving $m = .7$ also show an initial rapid rise in the response, followed by decline prior to the US which is determined by the value of h.

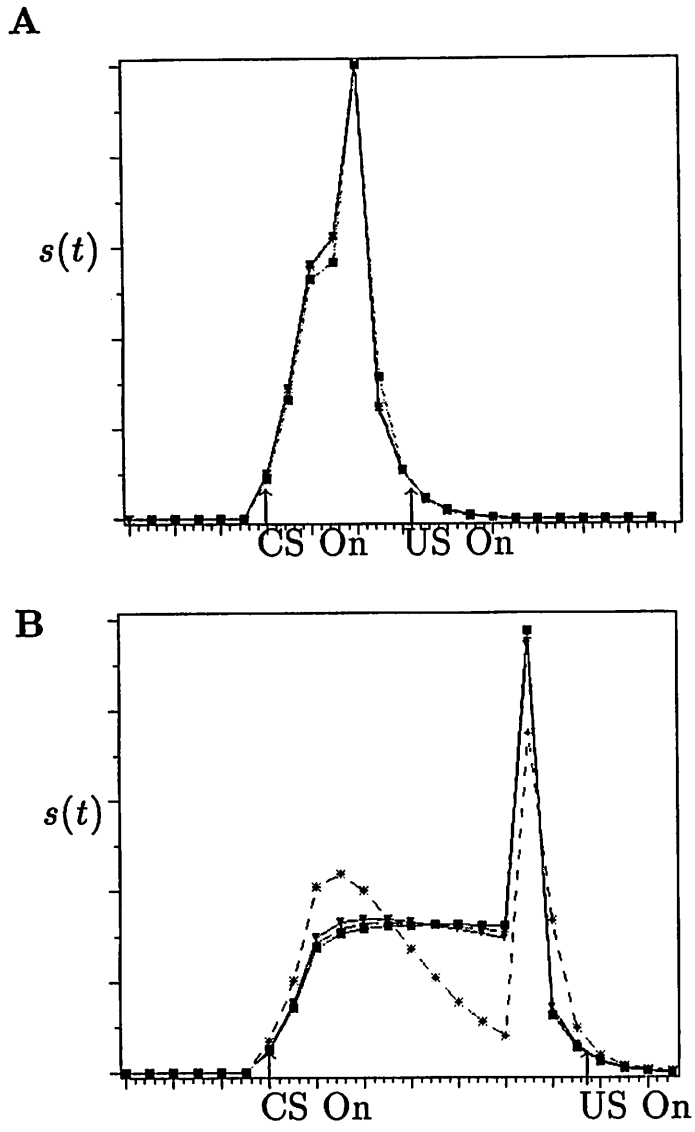

Figure D.7: **Response Topography as a Function of h.** Data shown were obtained at asymptote using a 50 ms US and ISI equal to CS duration. $\lambda = .7$. The starred plot is obtained with $h = .25$. In Panel A, CS duration = 300 ms. Values of Vs are as follows: with $h = 0.25$, V = .031; with $h = .75$, V= .103; with $h = 1.0$, V = .138; with $h = 2.0$, V = is .274. In Panel B, the CS duration is 650 ms. Values of V are as follows: for $m = 0.1$, $V = 2.2$; for $m = .35$, $V = .134$, and for $m = .7$, $V = .099$.

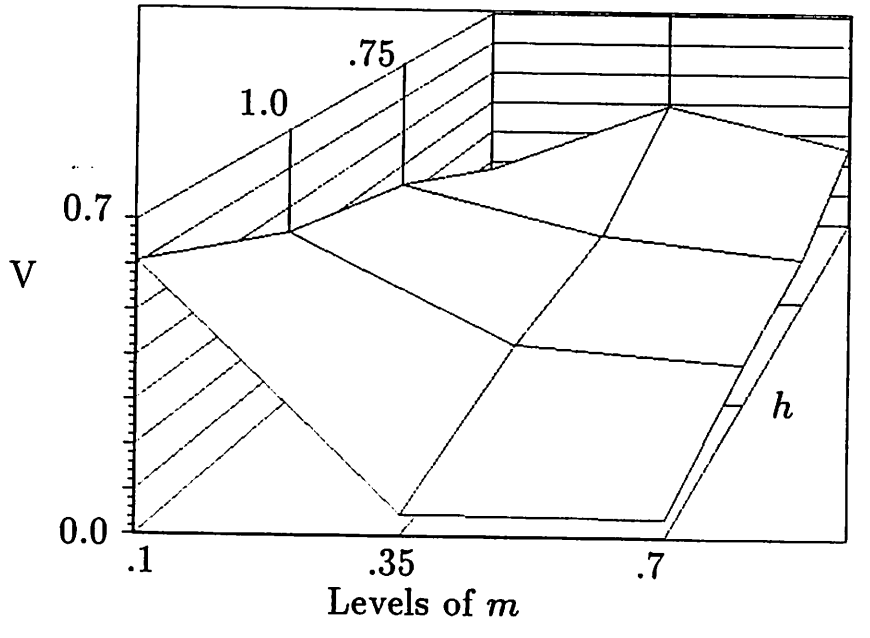

Figure D.8: V_i **as a Function of** m **and** h **in the Forward-delay Paradigm.** Data shown were obtained after 300 trials using a 50 ms US. Levels of parameter m are shown on the x-axis, levels of V_i are shown on the y-axis, with levels of parameter h shown on the remaining axis. Other parameter values were $\lambda = .7$; $\gamma = 0.9$; $k = 0.85$.

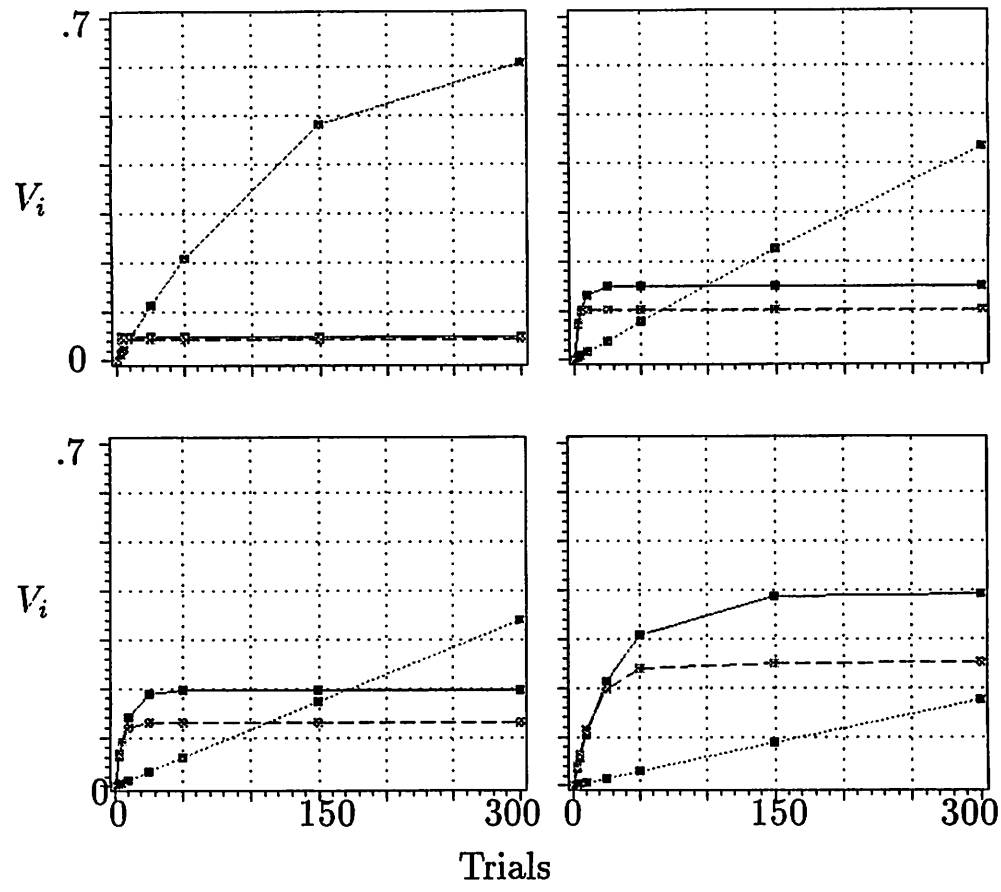

Figure D.9: **Acquisition of V_i as a Function of m and h, 300 ms CS.** Data shown were obtained after 300 trials using a 50 ms US and a 300 ms CS. The ISI was 300 ms. From upper left panel, moving clockwise, $h = .25$, $h = .75$, $h = 1.0$, $h = 2.0$. Other parameter values were $\lambda = .7$; $\gamma = 0.9$; $k = 0.85$.

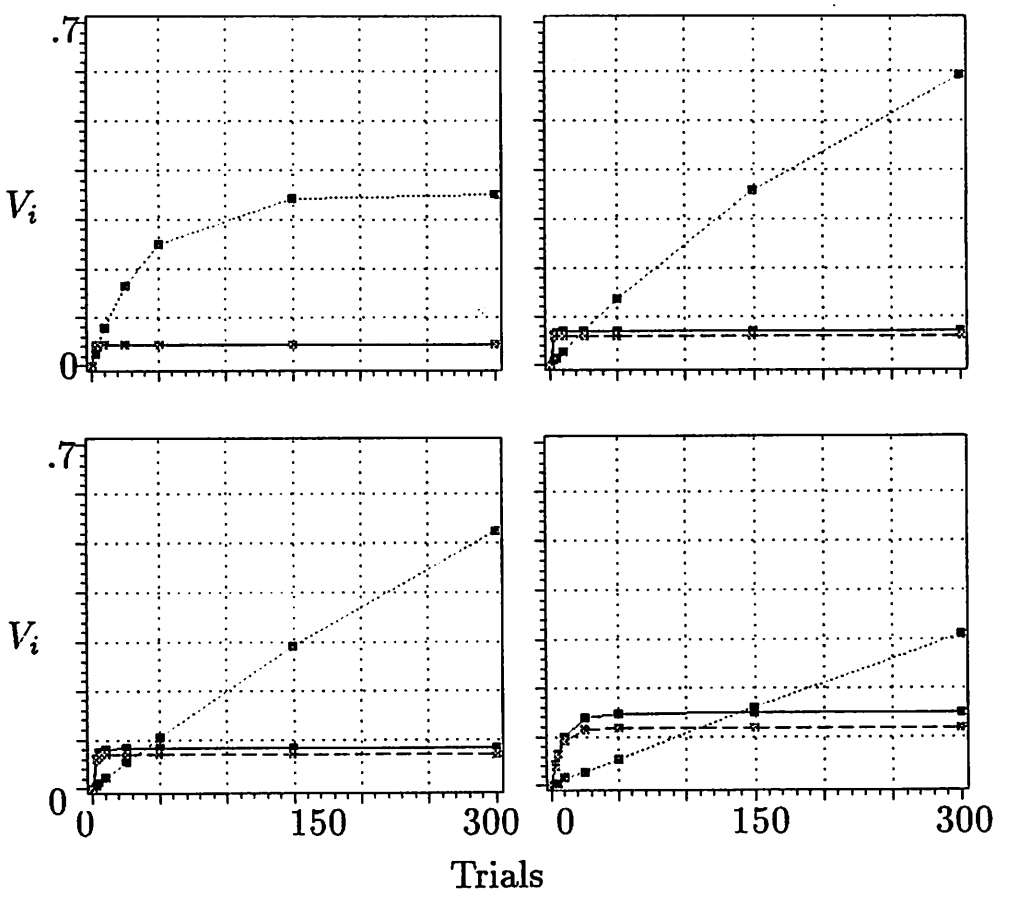

Figure D.10: **Acquisition of V_i as a Function of m and h, 650 ms CS.** Data shown were obtained after 300 trials using a 50 ms US and a 650 ms CS. The ISI was 650 ms. From upper left panel, moving clockwise, $h = .25$, $h = .75$, $h = 1.0$, $h = 2.0$. Other parameter values were $\lambda = .7$; $\gamma = 0.9$; $k = 0.85$.

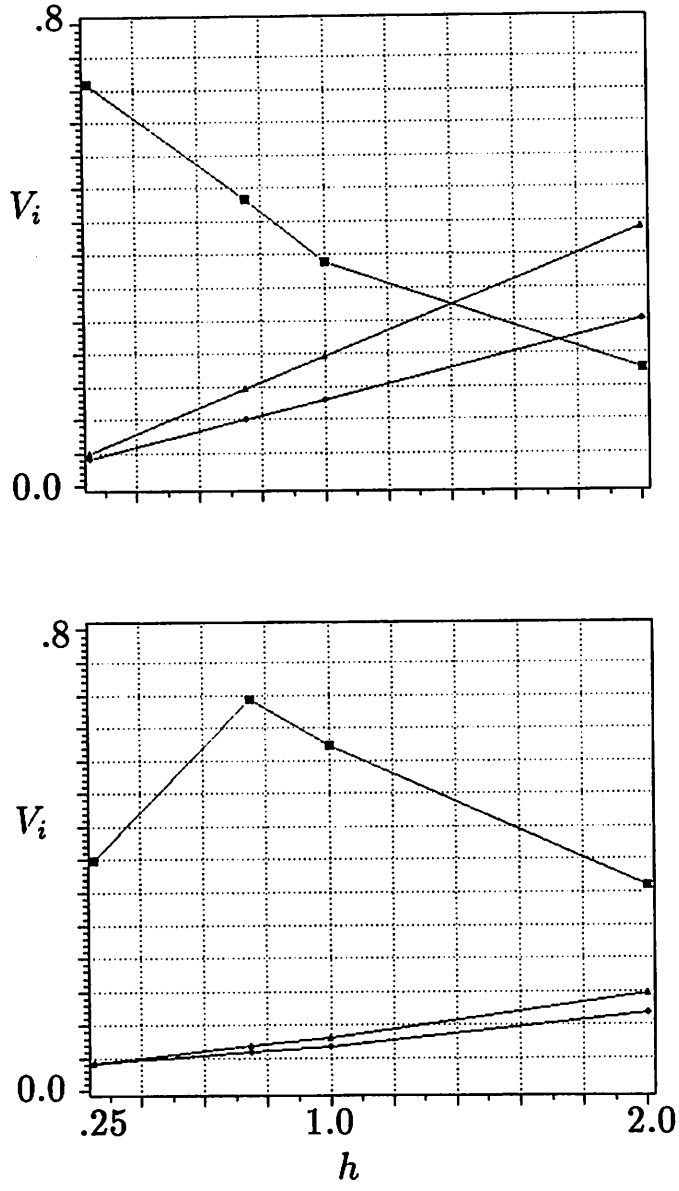

Figure D.11: **Acquisition of V_i as a Function of m, h, and ISI.** In A and B, ISI = CS Duration. From uppermost to lowest plots, $m = .1$, $m = .35$, and $m = .7$, respectively. In A, CS duration = 300 ms. In B, CS duration = 650 ms. Data shown were obtained after 300 trials using a 50 ms US. $\lambda = .7$; $\gamma = 0.9$; $k = 0.85$.

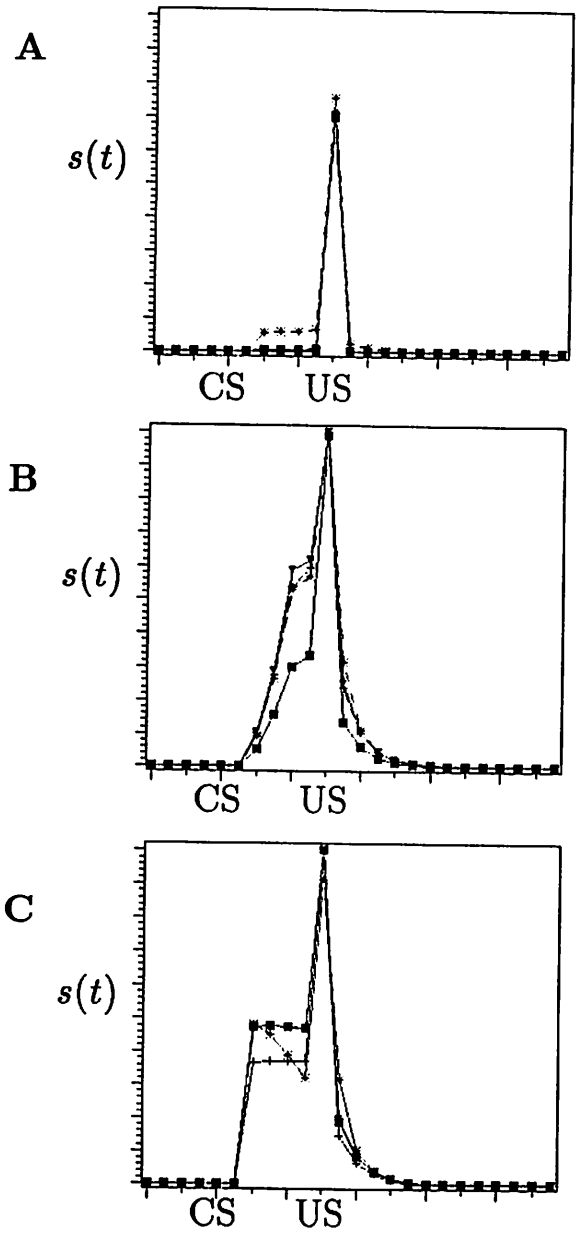

Figure D.12: **Response Topography as a Function of m and h, 300 ms CS.** Data were obtained using a 50 ms US, ISI = CS duration. CS and US labels designate stimulus onsets. Plots demarcated with stars indicate data obtained with $h = .25$, with triangles, $h = .75$, crosses, $h = 1.0$, and squares, $h = 2.0$. In A, $m = 0.1$. In B, $m = 0.3$. In C, $m = 0.7$. k is fixed at 0.85. 50 trials, 300 ms CS, 50 ms US, $\lambda = .7$.

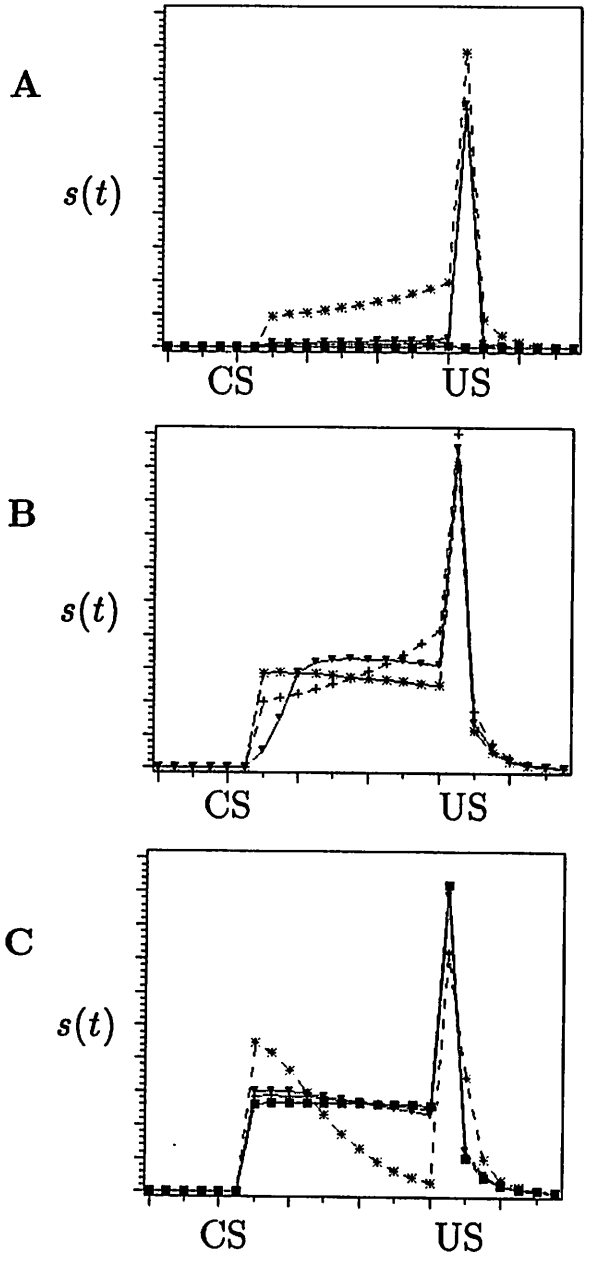

Figure D.13: **Response Topography as a Function of m and h, 650 ms CS.** Data were obtained using a 50 ms US, ISI = CS duration. CS and US labels designate stimulus onsets. Plots demarcated with stars indicate data obtained with $h = .25$, with triangles, $h = .75$, crosses, $h = 1.0$, and squares, $h = 2.0$. In A, $m = 0.1$. In B, $m = 0.3$. In C, $m = 0.7$. k is fixed at 0.85. 50 trials, 300 ms CS, 50 ms US, $\lambda = .7$.

APPENDIX E

SIMULATIONS: LEARNING RULE PARAMETERS

The simulation experiments in this section describe the effects of varying parameters of Equations 2.3 and 1.8 of the TD_{RTS} model. These parameters include γ, which determines the rate at which future values of λ are discounted, c, the learning rate parameter, and β, the eligibility trace recursion parameter. Acceptable parameter values should permit the TD_{RTS} model to describe ISI functions and reasonable topographies for optimal ISIs. Therefore, the experiments in this section focus on these key predictions. Parameters for the CS template x were set as follows: $m = 0.35$, $h = 1.0$, and $k = 0.85$.

Figure E.1, shows ISI functions obtained for values of γ ranging from .35 through .95. For these cases, the ISI function obtained after 100 trials maintains the characteristic concave-downward shape. With high γ, however, rate of acquisition at short ISIs is considerable slowed, thereby preventing the prediction of decreasing ISI functions after 100 trials. In simulations of the TD model, Sutton and Barto (1987) also used a high value for γ to control the values of V_i for short ISIs.

Response topographies produced by the TD_{RT} model are affected by the value of γ. Figure E.2 shows response topographies obtained on the 50th trial for a procedure with various values of γ and Equation 2.4. Panel A shows a 250 ms CS, a 50 ms US. The figure shows that CR amplitude decreases with increasing γ. For the TD_{RTS} model, V_i decreases with increasing γ, and by Equation 2.4, this leads to smaller CR amplitudes.

For the 500 ms CS shown in Panel B, CR amplitude tends to sag as γ increases, a consequence of the impact of γ on intra-trial V_i. Figure E.2 shows pronounced declines in CR amplitude that is minimized only when γ is greater than or equal to .85.

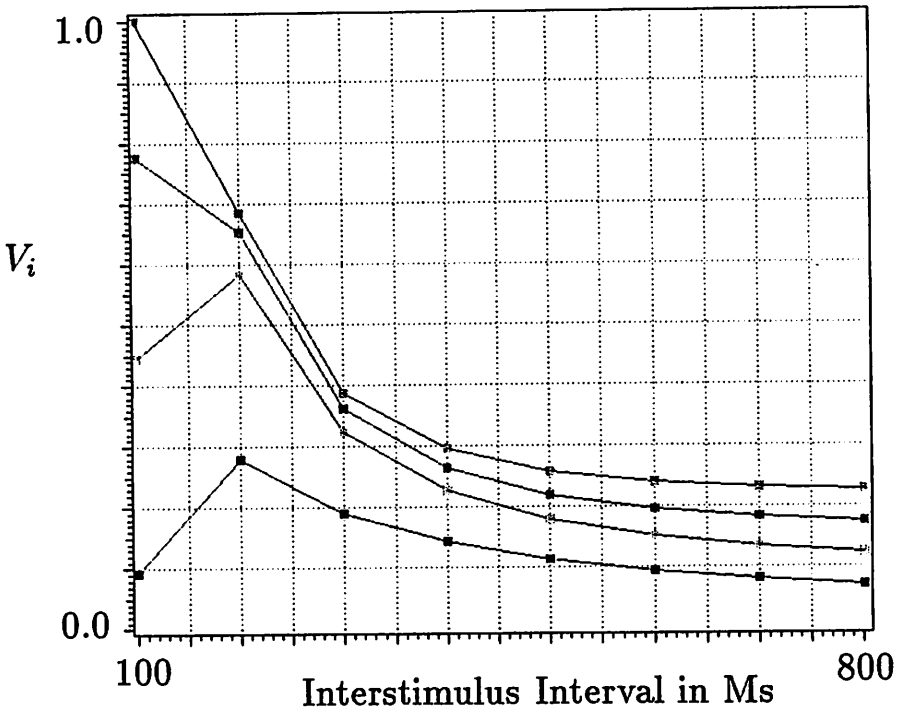

Figure E.1: V_i **after 100 Trials as a Function of ISI and Value of** γ. From upper to lower plots, parameter $\gamma = 0.35, .55, .75$, and .95. ISIs shown were forward-delay (CS duration = ISI) with a 50 ms US of height .5. Other parameter values included $c = 0.1$ and $\beta = .95$.

Figure E.3 shows an ISI function obtained after 50 trials for various values of β ranging from .25 to .95. As β increases, V_i for ISIs less than 250 ms decreases, but V_i for ISIs greater than 250 ms increases. In all cases the desired form of the ISI function is maintained.

Since parameter β has little effect on response topography, the results of experiments manipulating β are not shown. There is a tendency for CR amplitude to increase slightly with β for the 500 ms CS, owing to the fact that V_i increases with β at this ISI.

Generally, asymptotic values of V_i increase with c for both the TD and TD_{RTS} models. An explanation of the effect is shown with the TD_{RTS} model in Figure E.5, in which two values, $c = .1$ and $c = .5$, are compared. For a given starting value of V_i (in this case, .112) end-of-trial V_i increases with learning rate. As c increases, the value of V_i changes more during the course of the trial. Thus the unweighting generated by the presence of the CS is proportionally greater for a given starting value for V_i. At the time of the US, V_i is lower for the $c = .5$ case than for $c = .1$.

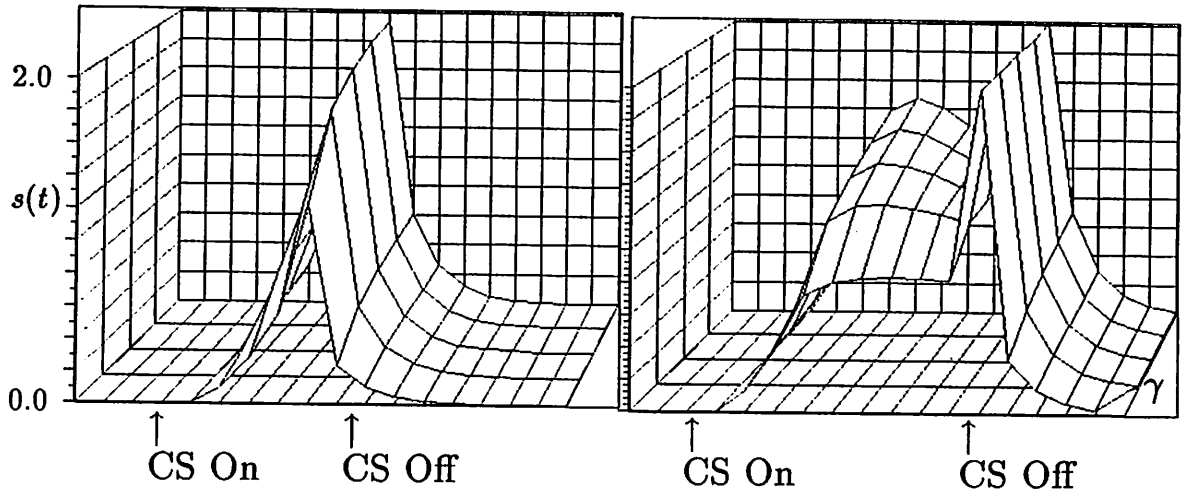

Figure E.2: **Response Topography after 50 Trials as a Function of Value of γ.** Parameter γ ranges from 0.55 to 0.95 (.95 closest to reader). Left Panel: The CS was 250 ms in duration; the US was of height .7 and 50 ms duration. Right Panel: The CS was 500 ms in duration. Other parameter values included $c = 0.1$ and $\beta = .95$. Output $s(t)$ is clipped at 2.0. In this and subsequent figures, topographies were computed using Equation 2.4 and a gain factor of 5.

Since V_i at CS offset is smaller for larger c, the amount of unweighting is smaller as well. Thus, by Equation 1.6 the net reinforcement generated during the US increases with c.

Figure E.4 shows the values of V_i obtained after 100 trials for three values of c. Figure E.4 shows that with a large value of c, for example $c = .5$, the characteristic inverted-U-shaped ISI function is not obtained. Instead, conditioning decreases with ISI for all the ISIs shown.

Parameter c exerts some control over response topography in the TD model, particularly at long ISIs. Figure E, Panel A shows that with a 250 ms CS, CR amplitude increases with c since V_i increases with c. In Figure E, Panel B, a 500 ms CS is depicted. Response topography is seen to drop off prior to the US to a greater degree as c increases. As we have already noted, V_i tends to decrease slightly over the CS duration due to unweighting caused by γ. Decrements in V_i are directly related to the value of c in Equation 1.6.

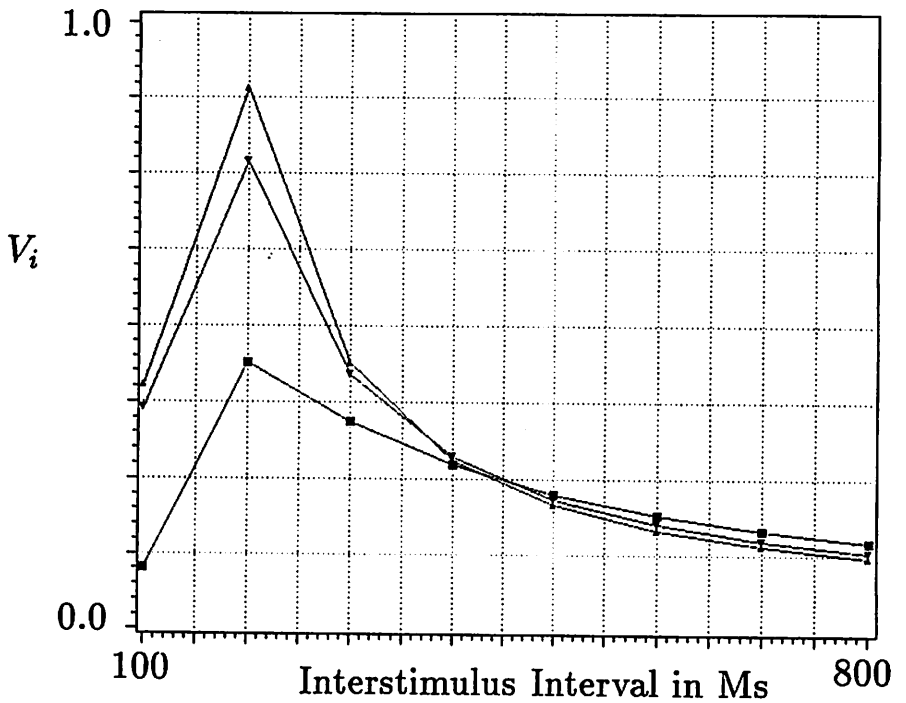

Figure E.3: V_i **after 100 Trials as a Function of ISI and Value of** β**.** Parameter β ranges from 0.25 (lowest plot) to 0.95 (uppermost plot). ISIs shown are for forward-delay conditioning with US onset immediately after CS termination; the US was of height .7 and 50 ms duration. Other parameter values included $c = 0.1$ and $\gamma = .95$.

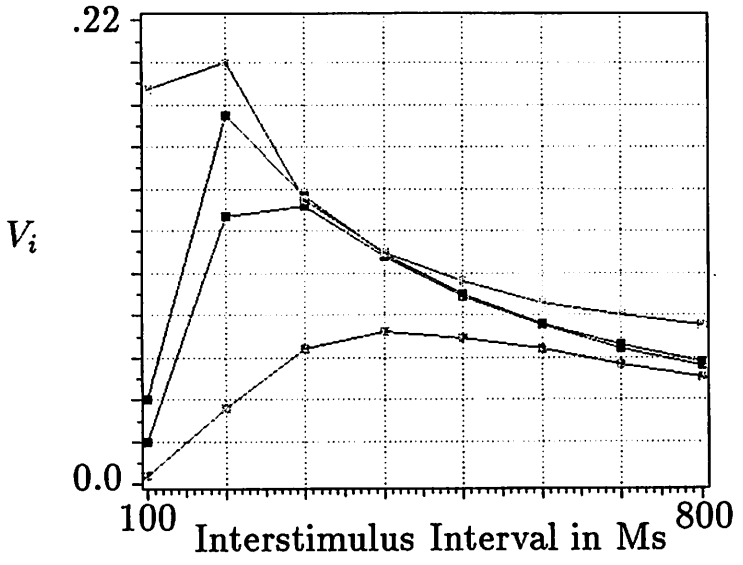

Figure E.4: V_i **after 100 Trials as a Function of ISI and Value of** c**.** Parameter c ranges from 0.01 (lowest plot) to 0.5 (uppermost plot). ISIs shown are for forward-delay conditioning with US onset immediately after CS termination; the US was of height .7 and 50 ms duration. Other parameter values included $\beta = \gamma = .95$.

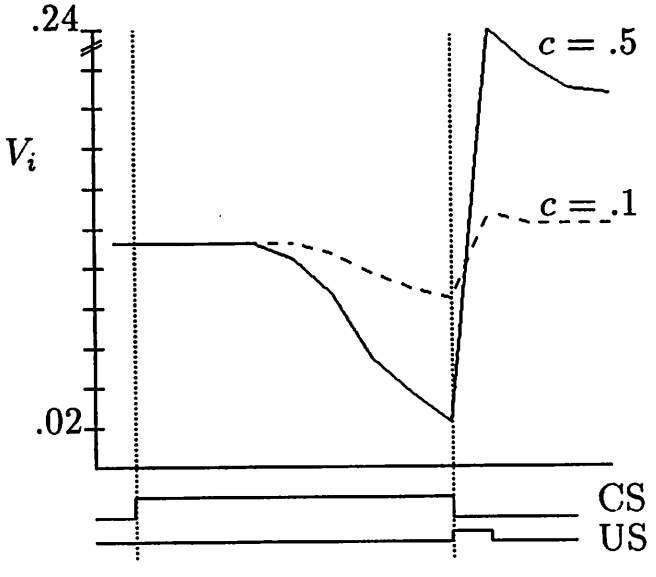

Figure E.5: **Change in Intratrial** V_i **as a Function of Value of** c**.** The CS is 400 ms; the US is 50 ms. V_i at the start of the trial is .11. Other parameter values include $\beta = 0.8$ and $\gamma = .8$.

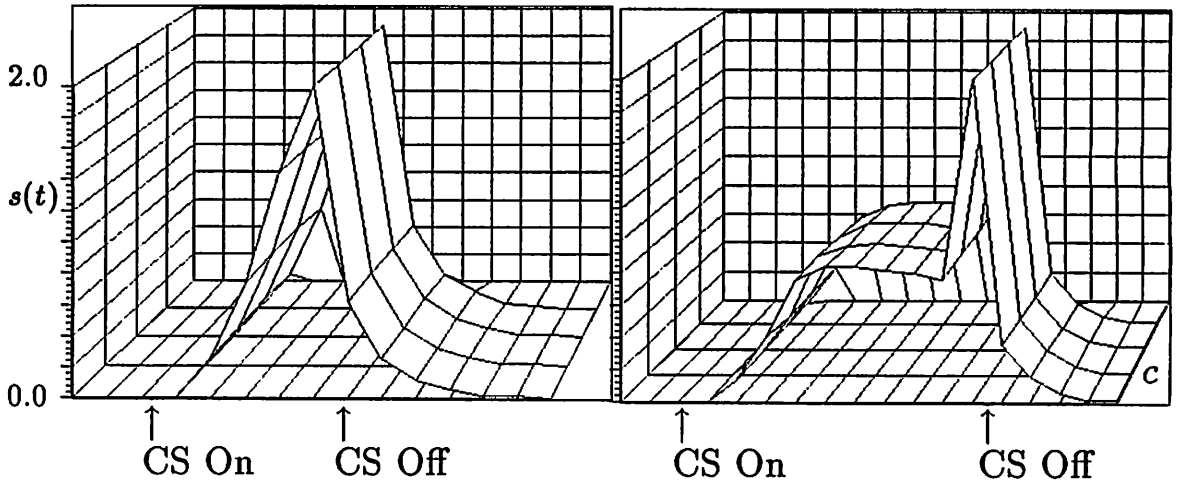

Figure E.6: **Response Topography after 50 Trials as a Function of Value of** c. Parameter c ranges from 0.01 to 0.5, with .5 closest to the reader. Panel A: The CS was 250 ms in duration; the US was of height .7 and 50 ms duration. Panel B: The CS was 500 ms in duration. Other parameter values included $\gamma = 0.95$ and $\beta = .95$.

APPENDIX F

FURTHER CONDITIONING SIMULATIONS

This appendix contains selected simulations of the TD_{RTS} model for a variety of conditioning protocols. The purpose of these experiments was to determine if the TD_{RTS} model required no further modification of its parameter values in order to account for phenomena beyond response topography and ISI functions. These phenomena, which are defined below, include backwards and simultaneous conditioning, US duration and intensity effects, trace conditioning, compound conditioning, conditioned inhibition, and blocking.

F.1 US Duration and Intensity Effects

Figure F.1 shows the value of V obtained after 50 trials using a 300 ms CS. The first aspect of the graph to consider is that conditioning is an increasing function of the value of λ, a prediction consistent with the findings of Smith (1968). Secondly, the TD_{RTS} model predicts that conditioning is a negatively-accelerated function of US duration. Tait, Kehoe, and Gormezano (1983) have shown that NMR conditioning is an increasing function of US duration.

Because eligibility, which determines the extent to which changes in weight occur, declines after CS offset, the TD_{RTS} model predicts that *there will be a US duration after which further increases will not result in additional associative strength.* For example, with a 300 ms CS, eligibility is zero about 500 ms after CS offset. Consistent with the theoretical result is the finding of Tait et al., (1983, Experiment 4). Tait and colleagues showed no differences in CR frequency between a group trained with a 1500 ms US and another trained with a 6000 ms US.

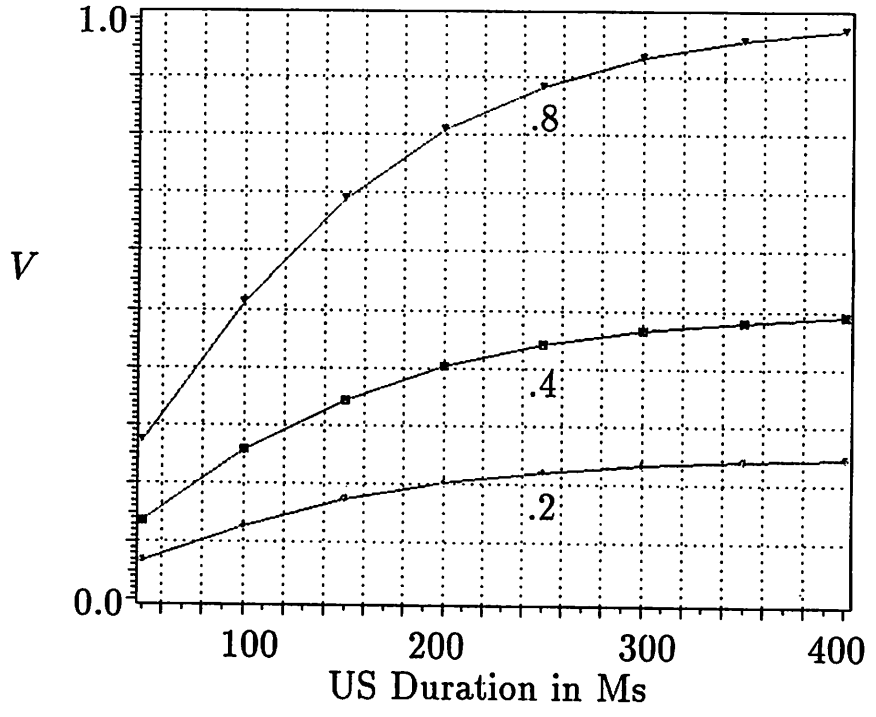

Figure F.1: V **As a Function of US Duration and Intensity.** V, shown on y-axis, was obtained after 50 trials with a 300 ms CS at an CS duration and ISI of 300 ms. US Intensity: .2 (lowest curve), .4 (middle curve), .8 (uppermost curve).

F.2 Conditioned Inhibition

Two trial types are involved in Pavlovian CI. One trial type pairs a CS (call it CSA) with the US; the second trial type pairs CSA with another CS (call it CSB); the compound is not reinforced. Over training, the animal learns to respond during the CSA+ trials but not the CSA-CSB trials. CSB has become a conditioned inhibitor; in the model this is represented as having a negative value of V. Figure F.2, Panel A illustrates that the TD_{RTS} model correctly predicts the results of such a procedure.

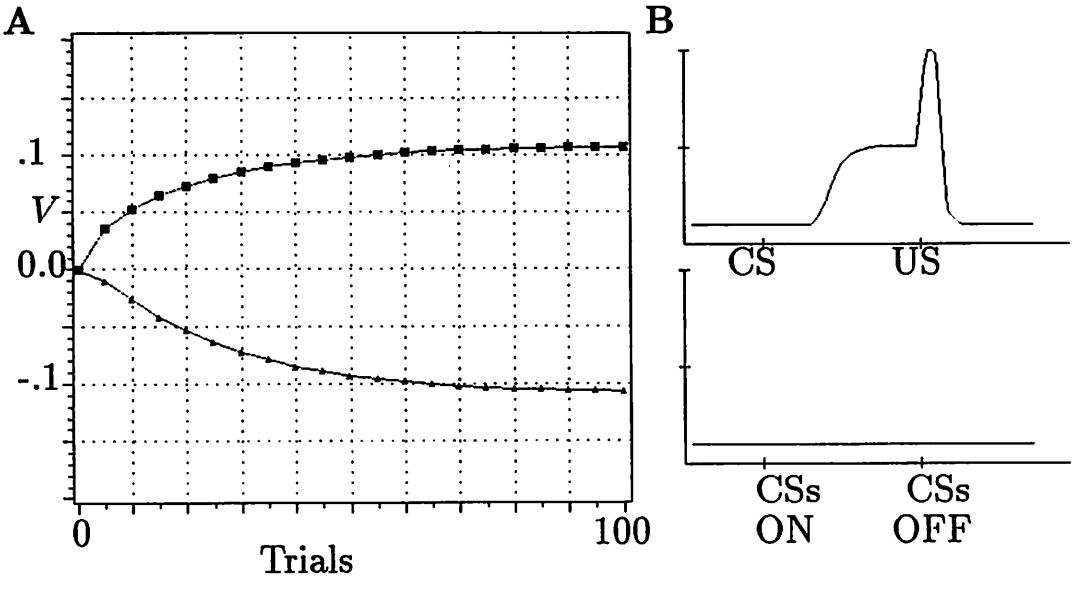

Figure F.2: **Acquisition During a Conditioned Inhibition Procedure.** Panel A: Acquisition of V over 100 trials for a conditioned inhibition procedure having two trial types: CS1, reinforced, alternated with a simultaneous unreinforced compound of CS1 and CS2. Durations of both CSs were 400 ms; the ISI on reinforced trials was 400 ms and the US duration was 50 ms. Panel B: Simulated Response profile after 25 trials for the trial type comprised of CS1 reinforced. Panel C: Response profile on trial 26, which is a presentation of the unreinforced compound. $\lambda = .7$.

Figure F.2, Panels B and C demonstrates the TD_{RTS} model's predictions for responding on a CI task; as observed experimentally, the model predicts CRs on CSA+ trials and no CRs on compound trials (Marchant, Mis and Moore, 1972; Marchant and Moore, 1974).

F.3 Serial and Simultaneous Compound Conditioning

Simulation studies examined the performance of the TD_{RTS} model in a number of compound conditioning procedures, such as blocking, higher-order conditioning, and sequential procedures that show primacy effects and facilitation of conditioning to CSs that occur long before the US. The procedures and results are summarized below.

Many learning theories are built upon the idea that successive events must occur closely together in time in order to become associated. Evidence from animal learning preparations suggests, however, that CSs which are quite remote from a US can become capable of initiating long sequences of behavior. For the rabbit NMR, the strength of conditioning decreases at longer ISIs and does not normally occur at all with ISIs of 2 to 4 seconds (Schneiderman, 1966; Schneiderman and Gormezano,1964). In rabbit NMR conditioning paradigms, the use of intervening or "bridging" stimuli has been show to greatly extend the ISI that can support conditioning (Kehoe, Gibbs,Garcia, and Gormezano, 1979). For example, studies by Kehoe et al., (1981) and Kehoe and Morrow (1984) showed enhanced conditioning to a CS1 presented in a serial compound with CS2 when compared to a CS1 trace control. Kehoe et al., (1981) were able to show that conditoning to CS1 could occur at at ISI of 18.75 seconds, an ISI well beyond the normal conditionable range for forward-delay or trace protocols. These types of studies are characterized as "remote associations" experiments.

The typical remote associations experiment with the rabbit NMR compares a compound consisting of two sequentially-occurring CSs with a trace control, consisting of the first CS (CS1) and a gap between the offset of the CS1 and the US which is equal to the duration of the second CS (CS2). Tests of the associative strength of each CS are conducted by giving non-reinforced probe presentations of CS1 to subjects which have been pretrained to either the Compound or Trace presentations. In the literature, a potentiation of conditioning to CS1 is inferred by an increased probability of responding to CS1 exhibited by a compound group when compared to trace controls.

Sutton and Barto (1987) demonstrated that the TD model can yield facilitation of remote associations. The present discussion extends their results by showing that at certain values of γ, the TD_{RTS} model predicts decrements of responding to stimuli remote from the US, relative to single CS controls.

Consider the following simulation experiment. In the Experimental condition, a 300 ms CS1 was followed immediately by a 300 ms CS2, and then by a 100 ms US of $\lambda = .9$. The Control condition consisted of presentations of the 300 ms CS1 followed by the US at an ISI of 600 ms. In both cases 100 trials were given. Parameters $\beta = .95$, $c = .1$, and γ was varied. With $\gamma = .8$, the TD_{RTS} model predicted that V_iCS1 in the Experimental condition was lower (-.992) than in the corresponding control (.034). With $\gamma = .95$, the associative strength for CS1 in the Experimental condition was no longer inhibitory but was lower than that obtained with the control. However, with $\gamma = .99$, simulation results were similiar to those obtained in the rabbit in that VCS1 was higher in the Experimental condition (.075) than in the Control (.034). It should be noted that the point at which a value of γ leads to model failure to predict the empirical results depends on the value of β, which determines the time course of eligibility. Were a lower value of β to be used, the possible range of γ values would be larger.

Experiments 3 and 4 of this dissertation examine the possibility of facilitation of CS1 when CS2 occurs at the same time as the US. The TD, SBD, and TD_{RTS} models all differ to some degree in predicting the results of this procedure. A complete description of the procedure and the models' results preface each experiment.

Second-Order Conditioning

In a second-order conditioning experiment, one CS can serve as a reinforcer for another CS. One second-order conditioning procedure involves a two-stage design. In Stage 1 of such an experiment, the first-order CS (CS1) is paired with a US, usually until a reasonably high level of conditioning has occurred. In Stage 2, a new CS (the second-order CS, CS2) is paired with CS1. Responding to CS2 increases initially, but as CS1 extinguishes, responding to both CSs eventually declines (Rashotte, Griffin, and Sisk, 1977).

The TD_{RTS} model replicates the finding of second-order conditioning results, but only as long as the value of γ is high. For example, one simulation experiment conducted used a 300 ms CS1 and CS2 at an ISI of 300 ms. The initial value of V_i for CS1 was .2. With $\gamma = .8$, after 20 trials levels of V_i were as follows: CS1 had dropped to .012, and CS2 had a value of −.085. Only when γ was high, that is, equal to .99, did the value of CS2 become excitatory. For the positive ISI used in this sample experiment, the prediction of conditioned inhibition accruing to CS2, as indicated by the development of a negatively-valued weight, is not supported by the experiments cited above. The loss of excitatory higher-order conditioning with certain values of γ provides another important constraint on the TD_{RTS} model.

For rabbit NMR conditioning, robust higher-order conditioning is obtained with a one-stage design that intersperses CS1-US and CS2-CS1 pairings (Kehoe, Feyer,and Moses, 1981). In so doing CS1 is protected from losing associative strength quickly. Empirically, CS1 tends to lose its capacity to elicit CRs if training is extended (Kehoe et al, 1981) a result that is interpreted as a buildup of inhibition. The SBD model readily predicts good second-order conditioning with the one-stage design, but again the TD_{RTS} model predicts weak second-order conditioning, and even conditioned inhibition if the value of γ is low (less than .95 with the 10 ms time step). Relevant simulations are considered in Appendix F.

Figure F.3 shows the value of V_i for the second-order CS CS2, as function of the ISI elapsing between it and CS1. The first-order CS, CS1, is pretrained to a weight of 0.4. ISIs shown range from −100 to +500.The figure shows inhibitory weights for CS1 if CS1 is preceded by or substantially overlaps with CS2. These inhibitory weights are consistent with studies of conditioned inhibition (Marchant and Moore, 1974), during which alternating presentations of CS1 + US and CS2 + CS1 (in simultaneous compound) result in inhibitory associations for CS1. As the CS2-CS1 ISI increases, the weight of CS2 increases to peak at an optimal ISI, and then declines. Unfortunately, second-order conditioning has been shown to occur at CS2-CS1 intervals of at least 2.4 seconds; the ISI function in Figure F.3 approaches 0 long before then.

Not shown in Figure F.3 are the results of the same simulation experiment using $\gamma < .99$. When γ is less than .95, CS2 becomes inhibitory at all of the ISIs simulated. This is because the amount of reinforcement generated by the previously conditioned CS1 is proportional to the value of γ. The higher the value of γ, the greater the amount of secondary reinforcement.

Blocking

As described in Chapter 2, the SBD, SB, TD, and TD_{RTS} models account for many aspects of behaviors obtained with the Kamin blocking paradigm (1968)[1]. For the TD_{RTS} model, blocking of conditioning to CSA (the added CS) is strongly influenced by the time courses of the stimuli in the compound. If CSB is fully trained and the compound is simultaneous, CSA acquires no associative strength, a result that replicates empirical observations (Kamin, 1968).

The TD_{RTS} model does not predict temporal primacy effects in blocking unless the value of γ is high. Figure F.4 shows the learning curves generated by the TD_{RTS}

[1]In blocking paradigms, one CS (CSB, the blocker) is trained to asymptote and then paired with another CS (CSA, for added). The compound of the novel and the pretrained CS is reinforced.

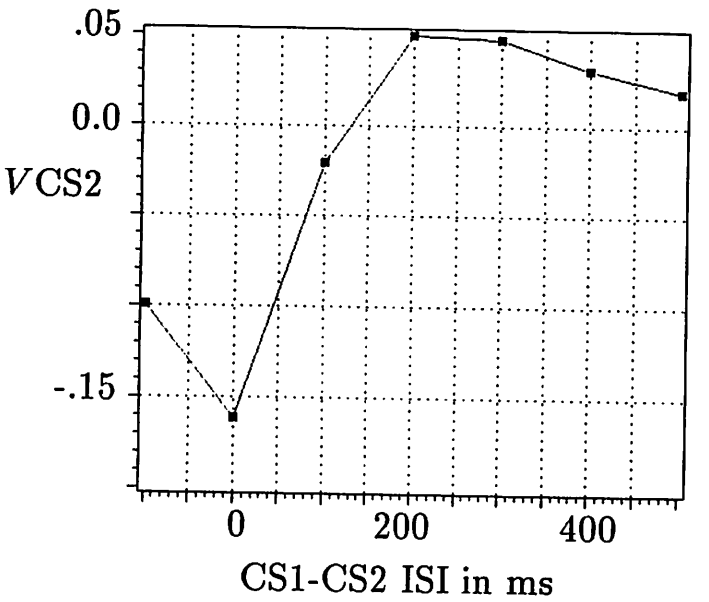

Figure F.3: **Second-Order Conditioning as a Function of the CS2-CS1 ISI.** CS1 and CS2 are each 300 ms in duration. V_i is obtained after 30 trials. The value of CS1 (the first-order CS) is .4 at the start of the simulation. Parameter $\gamma = .99$.

model when CSA begins *before* the onset of CSB. In Figure F.4, CSB loses some associative strength while CSA accrues strength, a result consistent with the findings of Kehoe et al., (1987). The simulation experiment shown in Figure F.4 used $\gamma = .99$. Not shown is the fact that the temporal primacy effect is also dependent on the value of γ. Were γ to equal .8, CSA would not gain strength.

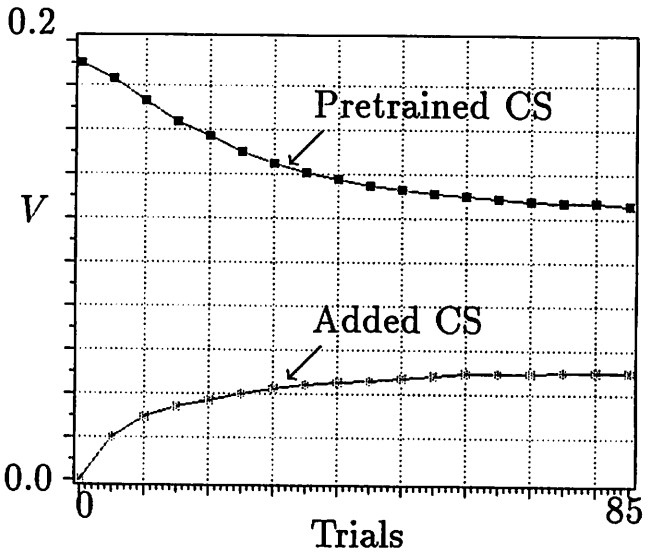

Figure F.4: **Effect of Temporal Primacy of the Added CS in Blocking.** Behavior of the model over trials in a blocking procedure in which the onset of the added CS occurs some time before that of pretrained CS. Initial weight of CSB was .19.

APPENDIX G

NMR DATA ANALYSIS SOFTWARE

At the University of Massachusetts conditioning laboratory, recording of the NMR is achieved via a linear optoelectronic phototransducer which communicates with the eye via a loop of suture sewn into the nictitating membrane; movement of the membrane produces a voltage change that is transmitted to an analog-to-digital converter installed in the Apple IIe computer. Following digitization, response waveform data are stored on floppy disk. These data are uploaded to a Sun Microsystems local area network where rapid analysis and high-resolution visualization are possible.

The data described in this dissertation were analyzed using NMRTOOL, software written in C to run under the Sun-OS operating system. NMRTOOL was designed to be consistent with data analysis programs written in FIRST and used by other laboratories that use the FIRST system for data acquisition and analysis. However, NMRTOOL provides rapid, detailed data analyses and better graphics than that obtainable with the Apple IIe.

NMRTOOL was developed by Diana Blazis with contributions from Bill Richards and John Desmond. It is an interactive program for analysis of data obtained during classical conditioning trials. NMRTOOL displays trial-by-trial waveforms and computes several response measures, including percentage CRs, peak and onset response latency, amplitude, and area. The program generates averaged waveforms (Figure G.1), plots learning curves, and distributions of response features. Hardcopy of NMRTOOL results is available at user request (Figure G.2).

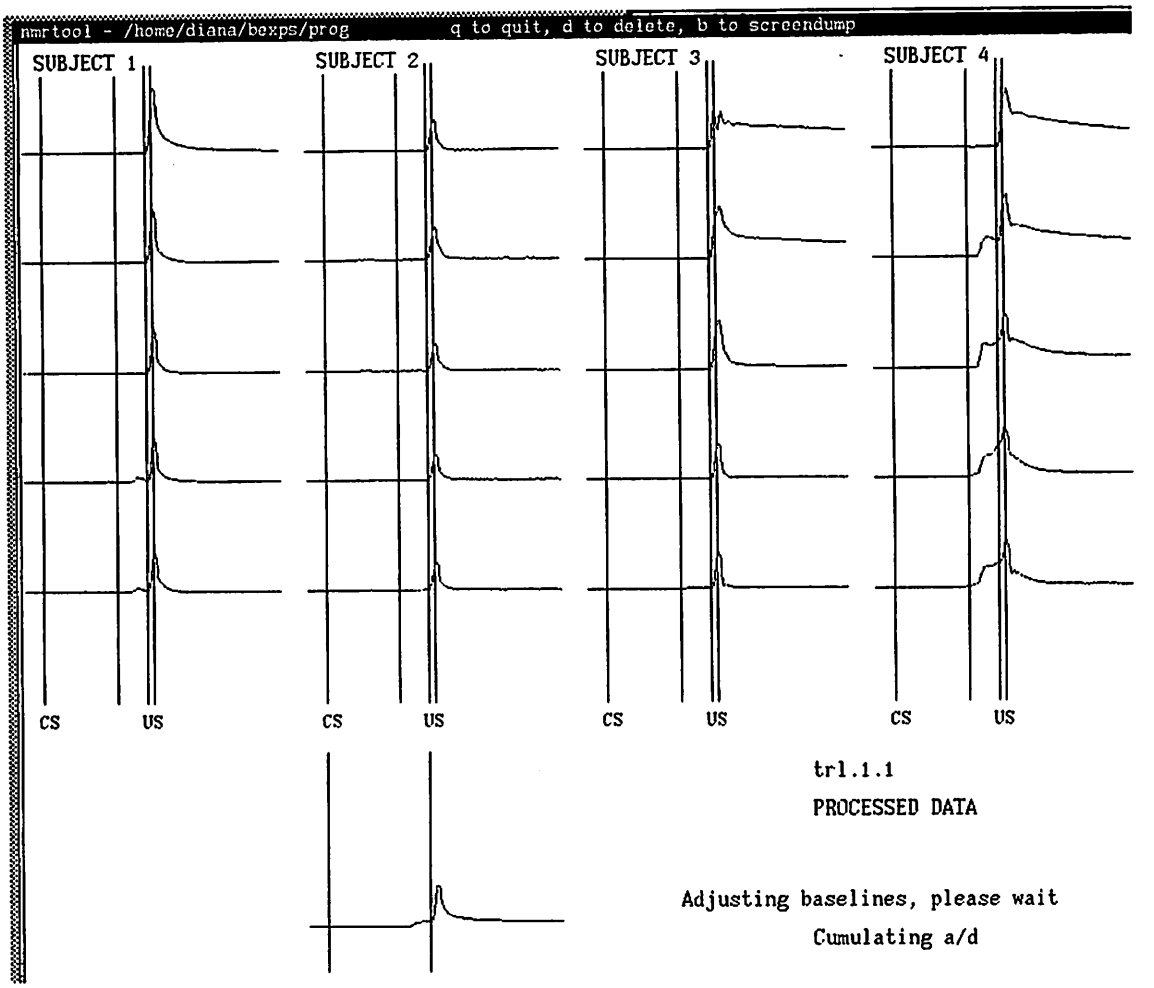

Figure G.1: **Output of NMRTOOL: Averaged Responses.** This screen of output of NMRTOOL shows responses averaged over 20 trials during a 100 trial session.

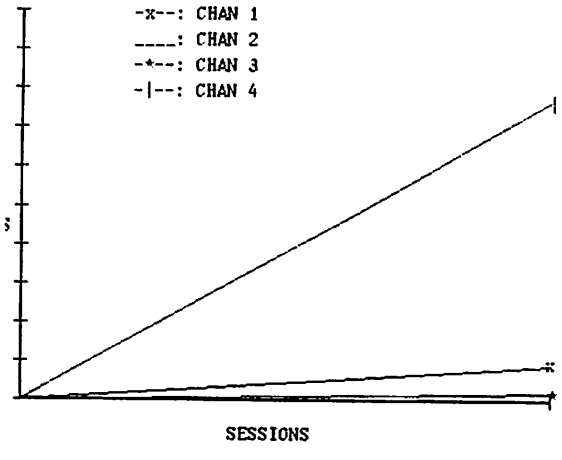

BLOCK TABLES	1	2	3	4	5	OVERALL
% CRs	0.0	0.0	0.0	15.0	30.0	9.0
ISI Amp., mm	0.06	0.07	0.09	0.36	0.38	0.2
ISI Area	16.8	5.1	5.5	20.3	38.1	17.1
UR Amplitude	3.15	2.58	2.13	2.22	1.98	2.4
UR Area	526.3	224.5	204.9	220.0	257.2	286.6
Onset Latency	819.5	942.0	995.0	870.5	804.5	886.3
Pk Lat (all TRL)	1027.5	1019.5	1015.0	1079.0	1010.5	1030.3
Pk Lat (CRs only)	489.5	248.0	458.0	870.0	789.0	570.9
% CRs	0.0	0.0	0.0	0.0	0.0	0.0
ISI Amp., mm	0.10	0.10	0.10	0.09	0.15	0.1
ISI Area	1.5	1.3	1.0	1.3	6.9	2.4
UR Amplitude	1.47	1.62	1.32	1.26	1.32	1.4
UR Area	191.4	230.6	153.5	149.4	170.9	179.2
Onset Latency	998.5	995.0	1013.5	999.0	959.0	993.0
Pk Lat (all TRL)	1038.5	1138.5	1037.0	1101.0	1066.0	1076.2
Pk Lat (CRs only)	289.5	341.5	349.5	377.5	567.5	385.1
% CRs	0.0	0.0	0.0	0.0	10.0	2.0
ISI Amp., mm	0.04	0.04	0.05	0.05	0.13	0.1
ISI Area	5.2	0.2	4.2	2.8	12.8	5.1
UR Amplitude	2.11	2.53	2.37	1.73	1.76	2.1
UR Area	1353.6	1313.5	417.2	190.9	199.2	694.9
Onset Latency	964.5	1018.5	952.0	971.0	977.5	976.7
Pk Lat (all TRL)	1061.5	1051.5	1040.0	1021.0	1061.0	1047.0
Pk Lat (CRs only)	589.5	395.0	387.5	318.5	420.5	422.2
% CRs	0.0	90.0	100.0	100.0	95.0	77.0
ISI Amp., mm	0.05	1.10	1.50	1.75	1.52	1.2
ISI Area	1.3	140.4	208.1	244.2	217.8	162.3
UR Amplitude	2.78	3.04	2.64	2.56	2.47	2.7
UR Area	1543.8	1528.5	1085.4	607.5	491.9	1051.4
Onset Latency	1017.5	839.5	792.5	761.0	789.5	840.0
Pk Lat (all TRL)	1081.5	1068.0	1048.5	1047.5	1043.5	1057.8
Pk Lat (CRs only)	668.0	901.5	915.5	943.5	912.0	868.1

Figure G.2: **Output of NMRTOOL: Topography Analyses and Percentage CRs.** This screen of output of NMRTOOL shows data analysis of response topography and an early learning curve.

Since NMRTOOL is approximately two thousand lines long, only a brief excerpt of the software is presented. The module shown is responsible for most of the analysis. Remaining code handles data and user I/O, graphics, some analysis, and reading and writing a data summary file.

```
/*   stat_nmr:   subroutine of NMRTOOL  */
/*   copyright Diana E.J. Blazis, 1989 */


/*_____________________________________INCLUDE_____________ */

#define NULLPR (struct pixrect *) NULL
#include <suntool/sunview.h>
#include <suntool/canvas.h>
#include <suntool/panel.h>
#include <math.h>
#include <stdio.h>
#include "nmr_def.h"
/*_____________________________________EXTERNS_____________ */
extern int addata[NTRLS][SUBJ][MAX_BINS_PER_TRIAL];
extern int pre_sel_flag[SUBJ][NTRLS]; /* preselected trial flags */
extern char filename[80]; /* data file to be processes */
extern int cslen[4], cs_start[4], cs_end[4],
us_start[4], us_end[4];  /* stimulus timers */

extern Frame frame; /* sunviews objects */
extern Canvas canvas;
extern Pixwin *pw;
extern struct pixfont *serif; /* serif font */
extern struct pixfont *lg_serif; /* 16 pt serif font */

extern int batchflag; /* flags for batch processing */
extern int stats_run;

/*_____________________________________________GLOBALS_________*/

int the_cr_amp[SUBJ][NTRLS],the_ur_area[SUBJ][NTRLS];
int the_cr_area[SUBJ][NTRLS], the_cr_lat[SUBJ][NTRLS]
int the_cr_flag[SUBJ][NTRLS], the_peak_lat[SUBJ][NTRLS];
int the_peak_lat_in_ISI[SUBJ][NTRLS],the_ur_amp[SUBJ][NTRLS];;


int accepted_trials[SUBJ][NTRLS]; /* array of acc trial flags */
short  block_cr_trials[SUBJ][BLOCK]; /* cr trials by block */
short block_acc_trials[SUBJ][BLOCK];/* accepted trials by block */
```

```
int  sub_acc[SUBJ]; /* acc trials by subject */
int  sub_crs[SUBJ]; /* num CRs by subject */

/* used for graphing crs over session */
static struct pr_pos vlist[4]={{100,400}, {100,396}, {104,400}, {104,396}};

void stat_canvas_process();

char *sprintf();
/*_____________________________________Main control__________ */

stat_nmr()
{


  find_cr_characteristics();
  window_set(canvas, WIN_CONSUME_KBD_EVENT, WIN_ASCII_EVENTS, 0);
  print_block_summaries();
  show_cr_distribution(pw);
  write_data_files(pw);
  write_dropped_trials();
  draw_lc();

  if(batchflag){
    system("screendump screentest");
    system("scr2ln3.new -z 2 -q -d screentest -l scr.dmp");
    system("lpr -h scr.dmp");
    system("sleep 30");
  }

  stats_run = 1;
  return;
}


/* _________________FIND CR CHARACTERISTICS_________________ */

/* calculate cr area, latency, amplitude, etc   */

find_cr_characteristics()
{
  extern int addata[NTRLS][SUBJ][MAX_BINS_PER_TRIAL]; /* digitized data */

  int peak_lat_in_ISI, peak_lat,cr_area, cr_lat;
  int ur_amp,  cr_amp, cr_flag, ur_area, latency;
  int i,j,k,l;
  int datum,  xbar; /* for rms calculations  */
  int addata_buf[BINS_PER_TRIAL];  /* local buffer */
```

```
  int latency_found,numacc;

  float x_bar,  squared_sum,sum ,numerator; /* also for rms */
  float term,sum_squared,denom; /* also for rms */
  double sqr_arg, rms;

  denom  = (float) ((cs_start[0]+5) * (cs_start[0]+4));
/* compute RMS during baseline to determine if trial
   should be tossed */

      sum_squared= 0.0;
      squared_sum = 0.0;
      sum = 0.0;
      for (j =0; j < (cs_start[0]+5); j++){
sum +=  (float) addata[i][k][j]; /* sum a/d data points */
/*  squared data points */
       sum_squared+=((float) addata[i][k][j]* (float) addata[i][k][j]);
      }
      x_bar = 0.0;
      xbar = 0;
      rms = 0.0;
      term = 0.0;
      squared_sum = sum * sum;
      sqr_arg = 0.0;
      numerator = ((cs_start[0] + 5) * sum_squared) - squared_sum;
      term += numerator/denom;
      sqr_arg = (double) term;
      rms = sqrt(sqr_arg);
      x_bar = sum/(cs_start[0] + 5);
      xbar = (int)(x_bar + .5);

bad_trl:
      if( rms > 1.0 )  /* bad rms trial */
{
      if(pre_sel_flag[k][i] == -1){
printf("\n Pre Sel Bad trial Sub %d Trial %d ", k+1, i+1);
      }
      else
{
  pre_sel_flag[k][i] == -1;
  printf("\n  Bad RMS trial Sub %d Trial %d ", k+1, i+1);
}
  ur_amp = -1;
  ur_area= -1;
  cr_amp = -1;
  cr_area = -1;
  latency = -1;
      cr_flag = -1;
```

```
      peak_lat = us_start[0]-cs_start[0];
      peak_lat_in_ISI = us_start[0]-cs_start[0];
  accepted_trials[k][i] = 0;
}
      else  /* rms ok */
{
  /* inits */
  ur_amp = 0;
  ur_area= 0;
  cr_amp = 0;
  cr_area = 0;
  latency = us_start[0] - cs_start[0];
  cr_flag = 0;
  accepted_trials[k][i] = 1;
  latency_found = 0;
  peak_lat=  us_start[0] -cs_start[0];
  peak_lat_in_ISI = us_start[0] -cs_start[0];

  /* baseline a/d  data */

  for(j = 0; j <= cs_start[0] + 4; j++){
 /* write analyzed data into (formerly) raw */
    addata_buf[j] = addata[i][k][j] - xbar + OFFSET;
    addata[i][k][j] = addata_buf[j];
  }

  /* ISI a/d  data */

  for( j = cs_start[0] + 5; j <= us_start[0]; j++){
    addata_buf[j]=  addata[i][k][j] - xbar + OFFSET;
            /* write analyzed data into (formerly) raw */
    addata[i][k][j] = addata_buf[j];
    if((addata_buf[j]-OFFSET) > cr_amp){
      cr_amp = addata_buf[j]-OFFSET;
      peak_lat = j -cs_start[0];
      peak_lat_in_ISI = j -cs_start[0];
    }

    if(latency_found ==0 && addata_buf[j] > OFFSET + 7) {
      found_latency_sub(j,&cr_flag,&latency,&cr_area, &ur_area,
addata_buf,&latency_found);
    }
    if(addata_buf[j] > OFFSET && latency_found == 1){
      cr_area += addata_buf[j] - OFFSET;
    }
  }

  for( j = (us_start[0]+1); j <= BINS_PER_TRIAL-1; j++){
```

```
    addata_buf[j]=  addata[i][k][j] - xbar + OFFSET;
    addata[i][k][j] = addata_buf[j];

    if((addata_buf[j]-OFFSET) > ur_amp){
      ur_amp = addata_buf[j]-OFFSET;
      if (ur_amp > cr_amp)peak_lat = j - cs_start[0];
    }
    /*  just offset here because it's possible to have really small URs */
    if(latency_found==0  && addata_buf[j] > OFFSET ) {
      found_latency_sub(j,&cr_flag,&latency,&cr_area,
&ur_area, addata_buf,&latency_found);
    }
    if (latency_found == 1 && addata_buf[j] > OFFSET)
ur_area += addata_buf[j]-OFFSET;
  }
}

      the_cr_area[k][i] = cr_area;
      the_cr_lat[k][i]  = latency;
      the_ur_amp[k][i] = ur_amp;
      the_ur_area[k][i] = ur_area;
      the_cr_amp[k][i] = cr_amp;
      the_cr_flag[k][i] = cr_flag;
      the_peak_lat[k][i] = peak_lat;
      the_peak_lat_in_ISI[k][i] = peak_lat_in_ISI;
    }


    }


/* cumulate cr trials and accepted trials per block */

  for (i=0 ;i < SUBJ; i++)   {
    sub_crs[i]= 0; /* init */
    sub_acc[i]= 0;
    for (j=0;j< BLOCK; j++){

      block_cr_trials[i][j] = 0;

      for(l=(j* TRL_PER_BLOCK); l < (j+1)*  TRL_PER_BLOCK; ++l) {

if (the_cr_flag[i][l] != -1 )
  block_cr_trials[i][j] += the_cr_flag[i][l];
block_acc_trials[i][j]+= accepted_trials[i][l];
if(the_cr_flag[i][l] != -1 )
  sub_crs[i] += the_cr_flag[i][l];
sub_acc[i] += accepted_trials[i][l];
```

```
      }
    }
  }
}


/* this routine counts back from when CR first exceeded threshold */

found_latency_sub(j,crflag,real_latency,crarea, urarea,
addata_buf,latencyfound)
     int j;
     int *urarea,*latencyfound;
     int *crarea,*crflag, *real_latency;
     int addata_buf[];

{

  int counter = j-1;

  *latencyfound = 1;

  if (j <= us_start[0] ) *crflag = 1;
/* find beginnning of response */
  while (addata_buf[counter] > OFFSET  ) {
    counter--;
  }
  *real_latency = counter+1;
  counter ++;
  while ( counter != j) {
    if (counter <= us_end[0]) *crarea  += addata_buf[counter]- OFFSET;
    if (counter > us_start[0]) *urarea += addata_buf[counter] - OFFSET;
    counter++;
  }

  *real_latency -= cs_start[0];

}
```

APPENDIX H

DATA FOR EXPERIMENT 2

Table H.1 provides a summary of data for all subjects of Experiment 2. The first column designates the Squad assignment. Squads denoted by "F" were those that received forward-delay conditioning; those denoted by "T" received trace conditioning. Squad names also contain a number, either 650 or 1000, which refers to the Squad's interstimulus interval in ms. Entries under "Session" are percentage CRs shown by a given subject on each daily session of training. Entries listed under "Trials to Criterion" are the number of trials each animal required to reach either the first CR, three CRs in a row, 5 CRs in a row, or 9 out of 10 CRs.

Table H.1. Percentage CRs over Days and Trials to Criterion, Experiment 2.

Squad	S	Session						Trials to Criterion			
		1	2	3	4	5	6	First CR	3 in Row	5 in Row	9 of 10
F650	1	0	.01	.54	.77	.87	.97	195	212	215	225
F650	2	.04	.11	.24	.75	.90	.78	80	246	248	339
F650	3	0	.24	.94	.52	.77	.92	149	173	206	210
F650	4	.62	.97	.95	.96	.90	.99	25	28	30	34
F1000	1	.01	.31	.64	.92	.89	.90	160	163	165	165
F1000	2	0	.46	.50	.76	.83	.99	108	157	180	191
F1000	3	0	0	.44	.49	.40	.76	216	219	221	278
F1000	4	.01	.06	.42	.19	.78	.90	111	135	196	199
T1000	1	.09	.06	.94	1.00	.98	.99	66	108	115	119
T1000	2	0	.69	.64	.82	.67	.97	102	105	110	115
T1000	3	.02	.71	.48	.99	.99	.96	83	123	156	160
T1000	4	.77	.88	.98	.98	.93	.97	21	24	26	30
T650	1	0	.40	.79	.89	.76	.90	128	151	153	164
T650	2	.70	.25	.52	.65	.74	.82	8	13	38	43
T650	3	.18	.11	.76	.87	.84	.98	44	73	75	210
T650	4	.62	.99	.65	.80	.75	.91	36	40	42	47

BIBLIOGRAPHY

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

Barto, A. G., & Sutton, R. S. (1982). Simulation of anticipatory responses in classical conditioning by a neuron-like adaptive element. *Behavioural Brain Research, 4*, 221-235.

Berger, T. W., Rinaldi, P. C., Weisz, D. J., & Thompson, R. F. (1983). Single-unit analysis of different hippocampal cell types during classical conditioning of the rabbit nictitating membrane response. *Journal of Neurophysiology, 50*, 1197-1219.

Berger, T. W., Weikart, C. L., Bassett, J. L., & Orr, W. B. (1986). Lesions of the retrosplenial cortex produce deficits in reversal learning of the rabbit nictitating membrane response: Implications for potential interactions between hippocampal and cerebellar brain systems. *Behavioral Neuroscience, 100*, 802-809.

Berthier, N. E. (1984). The role of the extraocular muscles in the rabbit nictitating membrane response: a re-examination. *Behavioural Brain Research, 14*, 81-84.

Berthier, N. E., & Moore, J. W. (1980). Role of extraocular muscles in the rabbit (*Oryctolagus cuniculus*) nictitating membrane response. *Physiology & Behavior, 24*, 931-937.

Blazis, D. E. J., & Moore, J. W. (1987). Simulation of a classically conditioned response: components of the input trace and a cerebellar neural network implementation of the Sutton-Barto-Desmond model. Department of Computer and Information Science, University of Massachusetts at Amherst, Amherst, Mass. Technical Report 87-74.

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 in rats. *Animal Learning and Behavior, 6*, 225-230.

Desmond, J. E. (1985). The classically conditioned nictitating membrane response: Analysis of learning-related single neurons of the brain stem. Ph.D. Dissertation. University of Massachusetts, Amherst, MA.

Desmond, J. E. (1988). Temporally adaptive conditioned responses: Representation of the stimulus trace in neural-network models. Computer and Information Science technical report 88-80, University of Massachusetts, Amherst, MA 01003.

Desmond, J. E., Blazis, D. E. J., Moore, J. W., & Berthier, N. E. (1986). Computer simulations of a classically conditioned response using neuron-like adaptive elements: Response topography. *Society for Neuroscience Abstracts, 12*, 516.

Desmond, J. E., & Moore, J. W. (1986). Dorsolateral pontine tegmentum and the classically conditioned nictitating membrane response: Analysis of CR-related activity. *Experimental Brain Research, 65*, 59-74.

Desmond, J. E., & Moore, J. W. (1988). Adaptive timing in neural networks: The conditioned response. *Biological Cybernetics, 58*, 405-415.

Donegan, N. H. (1981). Priming-produced facilitation or diminution of responding to a Pavlovian unconditioned stimulus. *Journal of Experimental Psychology: Animal Behavior Processes, 7*, 295-312.

Donegan, N. H., & Wagner, A. R. (1987). Conditioned diminution and facilitation of the UR: A sometimes opponent-process interpretation. In I. Gormezano, W. F. Prokasy, & R. F. Thompson (Eds.), *Classical conditioning, 3rd edition* (pp. 339-369). Hillsdale, NJ: Lawrence Erlbaum Associates.

Frey, P. W., & Sears, R. J. (1978). Model of conditioning incorporating the Rescorla-Wagner associative axiom, a dynamic attention process, and a catastrophe rule. *Psychological Review, 85*, 321-348.

Gabriel, M., Saltwick, S. E., & Miller, J. D. (1975). Conditioning and reversal of short-latency multiple-unit response in the rabbit medial geniculate nucleus. *Science, 189*, 1108-1109.

Gormezano, I. (1966). Classical conditioning. In J. B. Sidowski (Ed.), *Experimental Methods and Instrumentation in Psychology*. New York: McGraw-Hill.

Gormezano, I., Kehoe, E. J., & Marshall, B. S. (1983). Twenty years of classical conditioning with the rabbit. *Progress in Psychobiology and Physiological Psychology, 10*, 197-275.

Gormezano, I., Schneiderman, N., Deaux, E., & Fuentes, I. (1962). Nictitating membrane: Classical conditioning and extinction in the albino rabbit. *Science, 138*, 33-34.

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

Harvey, J. A., Gormezano, I., & Cool-Hauser, V. A. (1985). Relationship between heterosynaptic reflex facilitation and acquisition of the nictitating membrane response in control and scopolamine-injected rabbits. *The Journal of Neuroscience, 3*, 596-602.

Hayashi, H., Sumino, R., & Sessle, B. J. (1984). Functional organization of trigeminal subnucleus interpolaris: Nociceptive and innocuous afferent inputs, projections to thalamus, cerebellum, and spinal cord, and descending modulation from periaqueductal gray. *Journal of Neurophysiology, 51*, 890-905.

Hebb, D. O. (1949). *The Organization of Behavior.* New York: Wiley.

Hull, C. L. (1943). *Principles of Behavior.* New York: Appleton–Century–Crofts.

Hupka, R. B., Kwaterski, S. E., & Moore, J. W. (1970). Conditioned diminution of the UCR: Differences between the human eyeblink and the rabbit nictitating membrane response. *Journal of Experimental Psychology, 83*, 45-51.

Ison, J. R., & Leonard, D. W. (1971). Effects of auditory stimuli on the amplitude of the nictitating membrane reflex of the rabbit (Oryctolagus cuniculus). *Journal of Comparative and Physiological Psychology, 75*, 157-164.

Kamin, L. J. (1965). Temporal and intensity characteristics of the conditioned stimulus. In W. F. Prokasy (Ed.), *Classical conditioning* (pp. 118-147). New York: Appleton-Century-Crofts.

Kamin, L. J. (1968). Attention-like processes in classical conditioning. In W. F. Prokasy (Ed.), *Classical conditioning: A symposium* (pp. 118-147). New York: Appleton.

Kehoe, E. J. (1982). Conditioning with serial compound stimuli: Theoretical and empirical issues. *Experimental Animal Behavior, 1*, 30-65.

Kehoe, E. J., Feyer, A. M., & Moses, J. L. (1981). Second-order conditioning of the rabbit's nicitating membrane response as a function of the CS2-CS1 and CS1-US intervals. *Animal Learning & Behavior, 9*, 304-315.

Kehoe, E. J., Gibbs, C. M., Garcia, E., & Gormezano, I. (1979). Associative transfer and stimulus selection in classical conditioning of the rabbit's nictitating membrane response to serial conditioned stimuli. *Journal of Experimental Psychology: Animal Behavioral Processes, 5*, 1-18.

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

Kehoe, E. J., Schreurs, B. G., & Amodei, N. (1981). Blocking acquisition of the rabbit's nictitating membrane response to serial conditioned stimuli. *Learning and Motivation, 12*, 92-108.

Kehoe, E. J., Schreurs, B. G., & Graham, P. (1987). Temporal primacy overrides prior training in serial compound conditioning of the rabbit's nictitating membrane response. *Animal Learning and Behavior, 15*, 455-464.

Klopf, A. H. (1972). Brain function and adaptive systems– A heterostatic theory. Technical Rep. No. 133 [AFCRL-72-0164]. Air Force Cambridge Research Laboratories. L. G. Hanscom Field, Bedford, Massachusetts.

Klopf, A. H. (1988). A neuronal model of classical conditioning. *Psychobiology, 16*, 85-125.

Levinthal, C. F., Tartell, R. H., Margolin, C. M., & Fishman, H. (1985). The CS-US interval (ISI) function in rabbit nictitating membrane response conditioning with very long intertrial intervals. *Animal Learning and Behavior, 13*, 228-232.

Lipkin, S. G., & Moore, J. W. (1966). Eyelid trace conditioning, CS intensity, CS-US interval, and a correction for "spontaneous" blinking. *Journal of Experimental Psychology, 2*, 216-220.

Liu, S. S., & Moore, J. W. (1969). Differential conditioning of the rabbit nictitating membrane response: IV. Training based on stimulus offset and the effect of intertrial tone. *Psychonomic Science, 15*, 128-129.

Mackintosh, N. J. (1983). *Conditioning and associative learning.* New York: Oxford University Press.

Marchant, H. G., Mis, F. W., & Moore, J. W. (1972). Conditioned inhibition of the rabbit's nictitating membrane response. *Journal of Experimental Psychology, 95*, 408-411.

Marchant, H. G., & Moore, J. W. (1974). Below-zero conditioned inhibition of the rabbit's nictitating membrane response. *Journal of Experimental Psychology, 102*, 350-352.

Matzel, L. D., Held, F. P., & Miller, R. R. (1988). Information and expression of simultaneous and backward associations: implications for contiguity theory. *Learning and Motivation, 19*, 317-344.

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

Mis, F. W., Andrews, J. G., & Salafia, W. R. (1970). Conditioning of the rabbit nictitating membrane response: ISI X ITI interaction. *Psychonomic Science, 20*, 57-58.

Moore, J. W., & Blazis, D. E. J. (1989). Simulation of a classically conditioned response: A cerebellar neural network implementation of the Sutton-Barto-Desmond model.

In J. H. Byrne & W. O. Berry (Eds.), *Neural models of plasticity: Experimental and theoretical approaches* (pp. 187-207). New York: Academic Press.

Moore, J. W., & Desmond, J. E. (1982). Latency of the nictitating membrane response to periocular electrostimulation in unanesthetized rabbits. *Physiology & Behavior, 28*, 1041-1046.

Moore, J. W., Desmond, J. E., Berthier, N. E., Blazis, D. E. J., Sutton, R. S., & Barto, A. G. (1986). Simulation of the classically conditioned nictitating membrane response by a neuron-like adaptive element: Response topography, neuronal firing, and interstimulus intervals. *Behavioural Brain Research, 21*, 143-154.

Moore, J. W., & Stickney, K. J. (1980). Formation of attentional-associative networks in real time: Role of the hippocampus and implications for conditioning. *Physiological Psychology, 8*, 207-217.

Moore, J. W., & Stickney, K. J. (1985). Antiassociations: Conditioned inhibition in attentional-associative networks. . In R. R. Miller & N. E. Spear (Eds.), *Information processes in animals: conditioned inhibition* (pp. 209-222). Hillsdale, N.J: Lawrence Erlbaum Associates.

Pavlov, I. P. (1927). *Conditioned Reflexes (Translated by G.V. Anrep).* New York: Dover Publications, Inc.

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

Plotkin, H. C., & Oakley, D. A. (1975). Backward conditioning in the rabbit (*Oryctolagus cuniculus*). *Journal of Comparative and Physiological Psychology, 2*, 586-590.

Port, R. L., Romano, A. G., Steinmetz, J. E., Mikhail, A. A., & Patterson, M. M. (1986). Retention and acquisition of classical trace conditioned responses by hippocampal lesioned rabbits. *Behavioral Neuroscience, 100*, 745-752.

Rashotte, M. E., Griffin, R. W., & Sisk, C. L. (1977). Second-Order Conditioning of the Pigeon's Keypeck. *Animal Learning and Behavior, 5*, 25-38.

Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin, 72*, 77-94.

Rescorla, R. A. (1980). Simultaneous and successive associations in sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 6*, 207-216.

Rescorla, R. A. (1988). Behavioral studies of Pavlovian conditioning . In W. M. Cowan (Ed.), *Annual Review of Neuroscience* (pp. 329-352). Palo, Alto, California: Annual Reviews Inc.

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

Ryugo, D. K., & Weinberger, N. M. (1978). Differential plasticity of morphologically distinct neuron populations in the medial geniculate body of the cat during classical conditioning. *Behavioral Biology, 22,* 275-301.

Scandrett, J., & Gormezano, I. (1980). Microprocessor control and A/D acquisition in classical conditioning. *Behavior Research Methods & Instrumentation, 12,* 120-125.

Scavio, M. J., & Gormezano, I. (1974). CS intensity effects on rabbit nictitating membrane conditioning, extinction and generalization. *Pavlovian Journal of Biological Sciences, 9,* 25-34.

Schmajuk, N. A. (1986). Real-time attentional models for classical conditioning and the hippocampus. Ph.D. Dissertation, University of Massachusetts, Amherst, MA.

Schmajuk, N. A., & Moore, J. W. (1988). The hippocampus and the classically conditioned nictitating membrane response: A real-time attentional-associative model. *Psychobiology, 16,* 20-35.

Schmajuk, N. A., & Moore, J. W. (1989). Simulation of the classically conditioned nictitating membrane response by an attentional-associative network: Response topography, neuronal firing, and the effects of hippocampal lesions and stimulation. *Behavioural Brain Research, 32,* 173-189.

Schneiderman, N. (1966). Interstimulus interval function of the nictitating membrane response in the rabbit under delay versus trace conditioning. *Journal of Comparative and Physiological Psychology, 62,* 397-402.

Schneiderman, N., & Gormezano, I. (1964). Conditioning of the nictitating membrane of the rabbit as a function of CS–US interval. *Journal of Comparative and Physiological Psychology, 57,* 188-195.

Siegel, S. (1956). *Nonparametric Statistics for the Behavioral Sciences.* New York: McGraw-Hill Book Company.

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, 66,* 679-687.

Smith, M. C., Coleman, S. R., & Gormezano, I. (1969). Classical conditioning of the rabbit's nictitating membrane response at backward, simultaneous, and forward CS-US intervals. *Journal of Comparative and Physiological Psychology, 69,* 226-231.

Sutton, R. S. (1987). Learning to predict by the methods of temporal differences. Technical report TR87-509.1, GTE Lab, Waltham, MA.

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

Sutton, R. S., & Barto, A. G. (1987). A temporal-difference model of classical conditioning. Technical report TR87-509.2, GTE Lab, Waltham, MA.

Tait, R. W., Kehoe, E. J., & Gormezano, I. (1983). Effects of US duration on classical conditioning of the rabbit's nictitating membrane response. *Journal of Experimental Psychology: Animal Behavior Processes, 1*, 91-101.

Tait, R. W., & Saladin, M. E. (1986). Concurrent development of excitatory and inhibitory associations during backward conditioning. *Animal Learning and Behavior, 14*, 133-137.

Thompson, R. F., Donegan, N. H., Clark, G. A., Lavond, D. G., Lincoln, J. S., Madden, J., Mamounas, L. A., Mauk, M. D., & McCormick, D. A. (1987). Neuronal substrates of discrete conditioned reflexes, conditioned fear states, and their interactions in the rabbit. In I. Gormezano, W. F. Prokasy, & R. F. Thompson (Eds.), *Classical Conditioning III* (pp. 371-400). Hillsdale NJ: Lawrence Erlbaum Associates.

Weisz, D. J., & LoTurco, J. J. (1988). Reflex facilitation of the nictitating membrane response remains after cerebellar lesions . *Behavioral Neuroscience, 102*, 203-209.

Weisz, D. J., & Walts, C. (1990). Reflex facilitation of the rabbit nictitating membrane response by an auditory stimulus as a function of interstimulus interval. *Behavioral Neuroscience, 104*, 11-20.

Weisz, K. J., & McInerney, J. (1990). An associative process maintains reflex facilitation of the unconditioned nictitating membrane response during the early stages of training. *Behavioral Neuroscience, 104*, 21-27.

Yeo, C. H., Hardiman, M. J., & Glickstein, M. (1985a). Classical conditioning of the nictitating membrane response of the rabbit: I. Lesions of the cerebellar nuclei. *Experimental Brain Research, 60*, 87-98.

Yeo, C. H., Hardiman, M. J., & Glickstein, M. (1985b). Classical conditioning of the nictitating membrane response of the rabbit. II. Lesions of the cerebellar cortex. *Experimental Brain Research, 60*, 99-113.

Yeo, C. H., Hardiman, M. J., & Glickstein, M. (1985c). Classical conditioning of the nictitating membrane response of the rabbit: III. Connections of cerebellar lobule HVI. *Experimental Brain Research, 60*, 114-126.

Yeo, C. H., Hardiman, M. J., & Glickstein, M. (1986). Classical conditioning of the nictitating membrane response of the rabbit. IV. Lesions of the inferior olive. *Experimental Brain Research, 63*, 81-92.

Young, R. A., Cegavske, C. F., & Thompson, R. F. (1976). Tone-induced changes in excitability of abducens motoneurons and of the reflex path of nictitating membrane in rabbit *Oryctolagus cuniculus. Journal of Comparative and Physiological Psychology, 90*, 424-434.

Zimmer-Hart, C. L., & Rescorla, R. A. (1974). Extinction of Pavlovian conditioned inhibition. *Journal of Comparative and Physiological Psychology, 86*, 837-845.