Physiological Psychology
1980, Vol. 8 (2), 207-217

Formation of attentional-associative networks in real time: Role of the hippocampus and implications for conditioning

JOHN W. MOORE and KENNETH J. STICKNEY
University of Massachusetts—Amherst, Amherst, Massachusetts 01003

The function of the hippocampus in conditioning is portrayed in terms of an extension of Mackintosh's (1975) attention theory, which describes the evolution of the salience (associability) of each stimulus in the situation, including the context, and its predictive associative relationship to itself and all other stimuli. In terms of the model, the hippocampus is essential for computations that reduce salience when a stimulus is presented in the context of other stimuli that are better predictors of events. The model is applied to the phenomena of latent inhibition and blocking.

To be truly meaningful, any characterization of function, for example that of the hippocampal formation in learning and memory, must have implications for a set of variables, parameters, and transformations within a closed system of mathematical relationships. In science, the ability of a quantitative model to generate hard predictions, even wrong ones, accords it higher status than is typically accorded to purely verbal counterparts. Verbal theories are often vague, misunderstood, and misapplied.

O'Keefe and Nadel's (1978) bold proposal is that the role of the hippocampus in learning and memory can be aligned with the notion of the cognitive map. The hippocampus constructs such maps in order to use them in arriving at strategies (hypotheses) that guide behavior. These two concepts, that of the cognitive map and that of strategy selection, have been long abandoned by most learning theorists, not so much because these two cornerstones of S-S "expectancy" theory are wrong, but because these ideas have been vexedly difficult to portray in tractable mathematical terms. Part of the appeal of other hypotheses of hippocampal function, those expressed in terms of attention, inhibition, perseveration, orienting, etc., is that they appear to lend themselves more readily to a quantitative rendering.

Many contemporary theories of animal learning have retained the idea that S-S relationships or expectancies play a part in controlling an animal's behavior, but that they do so without assuming that cognitive maps are formed or that the animal engages in strategy selection. Most learning theorists use these concepts metaphorically. O'Keefe and Nadel's (1978) theory implies that they should be taken literally. In view of this challenge, it is incumbent on contemporary learning theory to begin developing an alternative framework for these ideas. Moreover, that alternative should take the form of a quantitative model.

Elsewhere, it has been suggested that the hippocampus is responsible for "tuning out" irrelevant events (Solomon & Moore, 1975). To date, this hypothesis has been as vague and open to misunderstanding as any of the other verbal theories, and its application to the phenomena of latent inhibition and blocking by Solomon and Moore (e.g., Moore, 1979a; Solomon, 1979) was based more on intuition than derivation. The purpose of this report is to present a formal mathematical model that predicts the phenomena of latent inhibition and blocking and also portrays the disruption of these effects in hippocampectomized animals by a common mechanism of "tuning out." No ready-made theories were available, but after consideration of three potential frameworks, the Rescorla-Wagner (1972) model, Wagner's (1976) processing model, and Mackintosh's (1975) attention theory, Mackintosh's model was selected as the most promising for further development. Mackintosh (1975) described the general conditions in which a conditioned stimulus (CS) controlling a learned response should gain salience and other instances in which salience should decrease. The experimental findings supporting our view of hippocampal function fit neatly within Mackintosh's account of salience loss. Solomon and Moore's (1975) "tuning out irrelevant stimuli" appeared to be synonymous with Mackintosh's (1975) "learning to ignore stimuli poorly correlated with reinforcement."

Preparation of this manuscript was facilitated by NSF Grant BNS 77-14871. Theoretical work was initiated in the spring of 1978, when the first author was on the staff of the MRC Unit on Neural Mechanisms of Behaviour, University College London, I. Steele-Russell, director. The authors are especially grateful to N. J. Mackintosh and A. Dickinson for their critiques of an earlier draft.

0090-5046/80/020207-11$01.35/0

Mackintosh's (1975) Theory

Mackintosh's (1975) theory revolves around two mathematical relationships, one describing the circumstances in which a target CS gains salience, the other describing the circumstances resulting in a loss of salience. Briefly, animals learn to attend to a target CS (A) to the extent that it uniquely predicts the outcome of a trial. The magnitude of the increase in the salience of A (α_A) is assumed to be proportional to the degree to which A is a better predictor of reinforcement than all other stimuli, including context, acting at the same time. Denoting "all other stimuli" by the symbol X,

$$\Delta\alpha_A > 0 \text{ if } |\lambda - V_A| < |\lambda - V_X|, \quad (1)$$

where V_A and V_X are the associative strengths of A and X with respect to the reinforcing event, and λ is the asymptotic limit of associative strength.

Associative values change according to familiar linear operators:

$$\Delta V_A = \alpha_A\theta(\lambda - V_A) \quad (2)$$

and

$$\Delta V_X = \alpha_X\theta(\lambda - V_X), \quad (3)$$

where $0 < \theta \leqslant 1$, $0 < \alpha_A$, $\alpha_X \leqslant 1$, and $\lambda > 0$. These equations (2-3) differ from the Rescorla-Wagner model in that, since strictly speaking A is compounded with X, changes in the associative value of A by the latter model are given by $\Delta V_A = \alpha_A\theta(\lambda - V_{AX})$, where $V_{AX} = V_A + V_X$.

Mackintosh (1975) accounts for reduction of α_A by the following relationship:

$$\Delta\alpha_A < 0 \text{ if } |\lambda - V_A| \geqslant |\lambda - V_X|. \quad (4)$$

The quantity $|\lambda - V_A|$ in Relationships 1 and 4 is the predictive "discrepancy" between the current associative value of A and the asymptotic value (λ) attainable through repeated application of Equation 2. Similarly, the quantity $|\lambda - V_X|$ is the discrepancy between "all other stimuli" and the maximum associative value implied by a particular reinforcer. Let $D_A = |\lambda - V_A|$ and $D_X = |\lambda - V_X|$. Relationship 1 states that α_A increases whenever D_A is smaller than D_X. Mackintosh (1975) states that the magnitude of this increase should be proportional to the difference between D_X and D_A.

Mackintosh's (1975) theory is asymmetrical in that no condition exists between D_A and D_X which permits α_A to remain constant (i.e., $\Delta\alpha_A = 0$). When A and X are equally good predictors of the trial outcome, both lose salience (Relationship 4). Symmetry between gains and losses of salience would be established by having $\Delta\alpha_A \leqslant 0$ replace $\Delta\alpha_A < 0$ in Relationship 4. In this case, α_A would remain constant as long as $D_A = D_X$; α_A would increase when $D_A > D_X$ and would decrease when $D_X > D_A$. If some rationale could be advanced to introduce symmetry into Mackintosh's (1975) theory, then it is a relatively simple matter to introduce quantitatively precise operators that capture the idea that changes (increments and decrements) of salience are proportional to the differential between D_A and D_X. Perhaps the simplest forms such operators can take are the following:

$$\Delta\alpha_A = C_A(1 - \alpha_A)(D_X - D_A)/2\lambda \quad (5)$$

whenever $D_A < D_X$, and

$$\Delta\alpha_A = -\alpha_A C_A(D_A - D_X)/2\lambda \quad (6)$$

whenever $D_A \geqslant D_X$; $0 < C_A < 1$.

The rate parameter C_A need not be equal in the two equations; for example, it might be easier for the target CS (A) to gain salience than to lose it, independent of the magnitude of the difference between D_A and D_X. Since the maximum possible difference between D_A and D_X is 2λ, for example, when $V_A \cong \lambda$ and $V_X \cong -\lambda$ (i.e., X is a conditioned inhibitor), the factor $1/2\lambda$ in Equations 5-6 ensures that α_A remains between 0 and 1. The factors $(1 - \alpha_A)$ and $-\alpha_A$ impose a "law of diminishing returns" on Equations 5-6 such that repeated application of Equation 5 results in a negatively accelerated approach of α_A to 1 from below, and repeated application of Equation 6 results in a positively accelerated approach of α_A to 0 from above. Equations 5-6, therefore, could replace Relationships 1 and 4 in any system of equations that describes the changes of value of all stimuli associated with reinforcement.

One gathers from his article that the asymmetry in Mackintosh's (1975) version of attention theory was necessitated by the phenomenon of latent inhibition. Latent inhibition refers to the retardation of conditioning that results from repeated preexposure of a CS prior to the initiation of paired presentations of the CS and the reinforcer. Most contemporary theorists concede that CS preexposure results in a loss of salience. The reduced value of α_A in Equation 2 implies slower learning than would be the case if the CS had not been preexposed.

In its present form, Mackintosh's Relationship 4 predicts latent inhibition by assuming that the CS (A) and context (X) are equally good predictors of trial outcomes (nonreinforcement) during the preexposure phase. Both A and X are assumed to have initial associative values of 0 (i.e., $V_A = V_X = 0$). Since it is customary to assume $\lambda = 0$ when stimulus presentations are not reinforced, Equations 2-3 are both

equal to 0. Therefore, $\Delta\alpha_A < 0$ because $|\lambda - V_A| = |\lambda - V_X|$ in Relationship 4. The target CS loses salience during each presentation (as does X) as the animal learns to ignore it. The accumulated loss must be made up during conditioning trials, and this handicap results in retarded acquisition by Equation 2.

Other Predictive Associations

The asymmetry of salience loss compared with salience gain in Mackintosh's (1975) model does not arise in Equations 5-6, where both gains and losses are proportional to the differential between D_A and D_X. Equations 5-6 might be substituted for Relationships 1 and 4 if latent inhibition can be derived from the original premises of attention theory in some other way. This becomes possible if other predictive associative relationships besides those involving the nominal reinforcer are taken into account.

Conditioned attention. Lubow and his associates have shown that if the to-be-conditioned stimulus (A) predicts some other stimulus of sufficiently high salience (B) during the preexposure phase of a latent inhibition experiment, then retarded acquisition to A in the second phase is reduced or eliminated (e.g., Lubow, Schnur, & Rifkin, 1976). For example, if a light is followed by a tone in a forward arrangement favorable for conditioning, subsequent pairings of the light with some reinforcer will not show retarded acquisition of the conditioned response (CR).

Although a formal representation of the foregoing paradigm is as yet unstated, Mackintosh's (1975) model can be applied quite easily by assuming that the salience of the target CS (α_A) depends on its relationship not only to the nominal reinforcer or unconditioned stimulus (US), but also to other stimuli as well. In Lubow et al.'s (1976) conditioned attention paradigm, the focal point of associative relationships prior to the introduction of the US is Stimulus B, the tone in our example. With notation and symbols altered heuristically, Relationships 1 and 4 become

$$\Delta\alpha_A > 0 \text{ if } |\lambda_B - V_A^B| < |\lambda_B - V_X^B| \tag{7}$$

and

$$\Delta\alpha_A < 0 \text{ if } |\lambda_B - V_A^B| \geq |\lambda_B - V_X^B|, \tag{8}$$

where

$$\Delta V_A^B = \alpha_A \theta_B (\lambda_B - V_A^B) \tag{9}$$

and

$$\Delta V_X^B = \alpha_X \theta_B (\lambda_B - V_X^B). \tag{10}$$

The quantities V_A^B and V_X^B are the associative values of A to predict B and X to predict B, respectively, and θ_B and λ_B are the rate and limit parameters that determine the associability of B with other stimuli.

If we let $D_A^B = |\lambda_B - V_A^B|$ and $D_X^B = |\lambda_B - V_X^B|$, the symmetrical form of attention theory can be substituted for Equations 7-8:

$$\Delta\alpha_A = C_A(1 - \alpha_A)(D_X^B - D_A^B)/2\lambda_B \tag{11}$$

when $D_A^B < D_X^B$ and

$$\Delta\alpha_A = -\alpha_A C_A (D_A^B - D_X^B)/2\lambda_B \tag{12}$$

when $D_A^B \geq D_X^B$.

We are suggesting that associations other than those revolving around the nominal reinforcing event determine the salience of the target CS, but that the rules that specify salience are basically the same in both instances. There are several reasonable hypotheses concerning the weight contributed by each stimulus in a given situation to the salience of the target. As an example, consider a situation like the one under discussion consisting of two noncompounded stimuli, A and B, a context X, and a nominal reinforcer (US). At any given moment of presentation, the salience of A could be derived from the weighted average of all the components of salience changes involving A in the associative network. The principal contributor would be the component $\Delta\alpha_A^{US}$, defined as follows:

$$\Delta\alpha_A^{US} = C_A(1 - \alpha_A)(D_X^{US} - D_A^{US})/2\lambda_{US} \tag{13}$$

when $D_A^{US} < D_X^{US}$, and

$$\Delta\alpha_A^{US} = -\alpha_A C_A (D_A^{US} - D_X^{US})/2\lambda_{US} \tag{14}$$

when $D_A^{US} \geq D_X^{US}$.

We have simply rewritten Equations 5-6 with superscripts to emphasize that $\Delta\alpha_A^{US}$ is to be regarded as but one component among potentially many similar relationships that govern $\Delta\alpha_A$. Another component in the present case would be $\Delta\alpha_A^B$ given by the right hand side of Equations 11-12.

Two sets of operators of identical form to Equations 11-14 must be added to complete the picture and to introduce an important idea, namely, that a stimulus can predict itself and thereby affect its own salience. This idea becomes important in developing a rationale for latent inhibition, which will be considered in the next subsection. One set of operators takes into account the association of A to X:

$$\Delta\alpha_A^X = C_A(1 - \alpha_A)(D_X^X - D_A^X)/2\lambda_X \tag{15}$$

when $D_A^X < D_X^X$, and

$$\Delta\alpha_A^X = -\alpha_A C_A (D_A^X - D_X^X)/2\lambda_X \tag{16}$$

when $D_A^X \geq D_X^X$.

Normally, A would lose salience based on its relationship with X because D_X^X is typically smaller than D_A^X. We are intentionally assuming that a stimulus (in this case X = context) can predict itself and that stable contextual stimuli are usually better predictors of themselves in any situation than are the more phasic occurrences of stimuli occurring within that context! The operators for V_A^X and V_X^X are of the same form as Equations 2-3 and 9-10.

$$\Delta V_A^X = \alpha_A \theta_X (\lambda_X - V_A^X) \tag{17}$$

and

$$\Delta V_X^X = \alpha_X \theta_X (\lambda_X - V_X^X). \tag{18}$$

Therefore,

$$D_A^X = |\lambda_X - V_A^X| \text{ and } D_X^X = |\lambda_X - V_X^X|.$$

Finally, since A can form a predictive relationship with itself such that one occurrence of A can predict the next, the component of salience change representing $\Delta\alpha_A^A$ is defined by the following set of equations.

$$\Delta\alpha_A^A = C_A(1 - \alpha_A)(D_X^A - D_A^A)/2\lambda_A, \tag{19}$$

when $D_A^A < D_X^A$,

$$\Delta\alpha_A^A = -\alpha_A C_A (D_A^A - D_X^A)/2\lambda_A, \tag{20}$$

when $D_A^A \geq D_X^A$,

$$\Delta V_A^A = \alpha_A \theta_A (\lambda_A - V_A^A), \tag{21}$$

and

$$\Delta V_X^A = \alpha_X \theta_A (\lambda_A - V_X^A), \tag{22}$$

where $D_X^A = |\lambda_A - V_X^A|$ and $D_A^A = |\lambda_A - V_A^A|$.

As we shall elaborate below, $\Delta\alpha_A^A$ will normally be negative because X is typically the better predictor of A than is A itself.

The change of salience of A in our example now becomes a weighted average of the following: $\Delta\alpha_A^{US}$, $\Delta\alpha_A^B$, $\Delta\alpha_A^X$, and $\Delta\alpha_A^A$.

How should these weights be selected and scaled? Following a suggestion of Mackintosh (personal communication), the components of salience should be scaled *in relation* to the nominal reinforcer in the situation, here the US. The weights for the superscripted stimuli are indicated as subscripts of the ϕs in the following equation

$$\Delta\alpha_A = \frac{\phi_{US}\Delta\alpha_A^{US} + \phi_B\Delta\alpha_A^B + \phi_X\Delta\alpha_A^X + \phi_A\Delta\alpha_A^A}{\phi_{US} + \phi_B + \phi_X + \phi_A}. \tag{23}$$

Since US would typically be represented in memory more strongly than any of the other stimuli, ϕ_{US} would ordinarily contribute the greatest weighting to $\Delta\alpha_A$. Indeed, most conditioning experiments are devised to ensure that the US is the dominant focal point of the associative relationships controlling an animal's behavior. In such circumstances, $\Delta\alpha_A$ is basically determined by $\Delta\alpha_A^{US}$, and the other contributing components might be disregarded with little loss of theoretical resolving power.

The denominator of Equation 23 increases whenever a new event occurs in context X, and once a given ϕ value has been established, it remains in the denominator for all subsequent computations of $\Delta\alpha_A$ even though that stimulus is not present during the point in time when the organism "computes" $\Delta\alpha_A$. In essence, we are assuming, for simplicity, that the organism has a perfect memory for events and can instantly assign a suitable weight to the event in Equation 23. The idea that animals can remember a single event is not too unreasonable (e.g., Blanchard, Shelton, & Blanchard, 1970), but the assumption that the magnitudes of the ϕs are set a priori defies our intuition of how memories are reevaluated in light of experience.

Latent inhibition. Prior to the introduction of the US, other associative relationships govern the salience of A. During its preexposure in a typical demonstration of latent inhibition, two associative relationships control the fate of α_A, that between A and X, and that between A and itself. The loss of salience of A during preexposure follows from the fact that X is typically a better predictor of A than is A itself. Accordingly, $\Delta\alpha_A^A < 0$ because $|\lambda_A - V_A^A| > |\lambda_A - V_X^A|$. Also, X is the better predictor of X than is A. Thus, $\Delta\alpha_A^X < 0$ because $|\lambda_X - V_A^X| > |\lambda_X - V_X^X|$, and Equation 23 becomes

$$\Delta\alpha_A = \frac{\phi_A\Delta\alpha_A^A + \phi_X\Delta\alpha_A^X}{\phi_A + \phi_X}.$$

Since the two terms in the numerator are negative, $\Delta\alpha_A$ is negative and salience is lost, yielding the latent inhibition (retardation) effect.

Under what circumstances might a stimulus such as A provide a better prediction of itself than that provided by the context in which it occurs? Whenever temporal conditioning is possible. Apart from contextual cues, which are always the best predictors of themselves in any situation, a punctate event such as the presentation of stimulus A is a good predictor of itself (i.e., its next occurrence) to the extent that temporal conditioning occurs. Temporal conditioning presumably follows the same parametric rules as other forms of conditioning in that the interstimulus interval (ISI) between successive occurrences of A

must be within a value capable of supporting conditioning. In a rabbit eyeblink conditioning experiment, the optimal ISI is between .1 and .5 sec (Gormezano & Moore, 1969). Conditioning can occur at long ISIs up to 3 sec, but only when the CS and US are relatively intense.

A typical demonstration of latent inhibition in the rabbit nictitating membrane response (NMR) preparation (e.g., Solomon & Moore, 1975) entails presenting the to-be-conditioned CS, usually a pure tone of 75 dB (SPL) of .5 sec duration, once every 30 sec for 100 presentations per day and for 4.5 daily sessions prior to introduction of the eyeshock US. Our intuition tells us that it might prove difficult to obtain a latent inhibition effect if the parameters of the preexposure phase were changed to permit temporal conditioning of the tone. Temporal conditioning in this case obviously does not imply any peripherally observable conditioned response, only that the animal comes to think that one occurrence of the tone implies another. Suppose, therefore, that the 450 preexposures of Solomon and Moore (1975) were compressed into 7.5 min with a constant ISI of 1 sec (from onset to onset) rather than 30 sec. This procedure should lead to an increase, rather than a progressive decrease, in the tone's salience to the extent that $D_A^A < D_X^A$, i.e.,

$$A\alpha_A^A = C_A(1-\alpha_A)(D_X^A - D_A^A)/2\lambda_A.$$

If the memorial representation of the tone (ϕ_A) is considerably greater than that of the context (ϕ_X), Equation 23 yields

$$\Delta\alpha_A \cong \Delta\alpha_A^A.$$

The foregoing analysis leads to the prediction that latent inhibition will not be observed. We choose to interpret this hypothetical experiment in terms of temporal conditioning. Others might prefer to think in terms of "sensitization" (cf. Groves & Thompson, 1970). Either way, A may not lose salience to the extent that it would if each occurrence were better predicted by context than by A itself.

Rates of presentation that are too fast, for example, with ISIs less than .1 sec from onset to onset of a stimulus of very short duration, would not necessarily result in temporal conditioning in the sense of the model, but to configuring instead. The entire train of pulsed stimulus presentations would likely be perceived as a qualitatively different stimulus, and no predictions of the model would be possible regarding the salience of individual pulses in the train.

Attentional-Associative Networks

In general, given a set of r stimulus, A_i, i = 1, . . . , r, with saliences α_i in context X, the change in salience of A_i is given by

$$\Delta\alpha_i = \frac{\sum_{k=1}^{r} \phi_k \Delta\alpha_i^k + \phi_X \alpha_i^X}{\sum_{k=1}^{r} \phi_k + \phi_X}. \quad (24)$$

Equation 24 is simply an extension of Equation 23. The superscripts designate the relevant discrepancies from a system of difference equations of the following form:

$$\Delta V_i^k = \alpha_i \theta_k (\lambda_k - V_i^k), \quad (25)$$

i, k = 1, . . . , r, and X,

$$\Delta\alpha_i^k = C_i(1-\alpha_i)(D_X^k - D_i^k)/2\lambda_k, \quad (26)$$

when $D_i^k < D_X^k$,

and

$$\Delta\alpha_i^k = -\alpha_i C_i (D_i^k - D_X^k)/2\lambda_k, \quad (27)$$

when $D_i^k \geq D_X^k$.

More Than Two Discrepancies—Compounding

When stimuli are compounded, for example, when A_i and A_k are presented together in context X, the change of α_i depends primarily on the predictive relationship of A_i to an arbitrarily selected reinforcing event. If we denote this reinforcing event US, and if we assume that it dominates the other elements of the attentional-associative network that determines the changes in α_i, then

$$\Delta\alpha_i \cong \frac{\phi_{US}\Delta\alpha_i^{US}}{\phi_{US}} \cong \Delta\alpha_i^{US},$$

as described above.

The problem that arises when A_i and A_k are compounded is in choosing the appropriate discrepancy for computation of $\Delta\alpha_i$. Either $\Delta\alpha_i$ can be computed from Equations 13-14 on the basis of the differential between D_i^{US} and D_k^{US} (the discrepancy of the other element of the compound) or on the basis of the differential between D_i^{US} and D_X^{US} (the discrepancy of the context). What is needed is a rule for choosing the appropriate discrepancy to represent "all other stimuli" for comparison with the discrepancy of the target stimulus. Of several reasonable approaches, the one adopted here is simply to choose whichever is smaller.

Application to Real Time

As portrayed thus far, the model tacitly assumes that the predictive associative relationship V_A^B between

two events, A and B, increases only when they occur simultaneously. In truth, empirical relationships between the rate of conditioning and the interstimulus interval (ISI) between the onset of A and the onset of B indicate that simultaneity yields weak conditioning at best. Optimal conditioning occurs when A precedes B by a short interval. Rate of conditioning falls off progressively as the interval of time (Δt) between A and B increases.

These facts can be taken into account by introducing an additional parameter in the equation for ΔV_A^B. This parameter, τ, is a function of Δt, the interval between A and B.

$$\tau = \tau(\Delta t) = e^{k(q-\Delta t)} \tag{28}$$

when $\Delta t \geqslant 0$. The constant q is positive and equal to the optimal ISI for conditioning; $k > 0$ is the rate constant. For purposes of computer simulations described below, τ was assigned an arbitrary small fraction value ($\tau = s$) for instances where $0 \leqslant \Delta t < q$. For cases where $\Delta t < 0$ (backward conditioning), τ was assigned the value zero. In the simulations of the model described below, the rate constant k was always .67 and $\tau = s = .067$.

Incorporating the ISI parameter into the model yields the following expression

$$\Delta V_A^B = \alpha_A \theta_B \tau(\Delta t)(\lambda_B - V_A^B).$$

When the interval between A and B is optimal ($\Delta t = q$) and A precedes B, Equation 29 reduces to the form described previously.

$$\Delta V_A^B = \alpha_A \theta_B (\lambda_B - V_A^B).$$

As the interval between A and B increases, ΔV_A^B progressively decreases by the fractional amount given by Equation 28. Whenever A and B occur simultaneously, or at a less-than-optimal separation, ΔV_A^B is a small fraction of its value when $\Delta t = q$.

In our simulations, there are occasions when A occurs more than once before B. This suggests that both proximal and remote associative changes could be computed simultaneously. For example, in three successive units of time of equal duration (bins), the sequence AAB implies that at the time B occurs, ΔV_A^B should be calculated with respect to A in the first bin and also with respect to A in the second bin. Our simulation selected the maximum absolute value of ΔV_A^B and discarded the other. Whether the calculation chosen would be the one based on A in the first bin depends on the size of the bin in relation to q, the optimal ISI.

Nonoccurrences of Events

So far we have not considered the ways in which nonoccurrences of previously experienced events influence the system. It seems reasonable to assume that the nonoccurrence of the event weakens its associative relationships with respect to its predictors. The size of the decrements in V_A^B occasioned by A predicting the nonoccurrence of B (non-B) is given by the following "extinction" operator:

$$\Delta V_A^B = \alpha_A \theta'_B \tau(\Delta t)(0 - V_A^B). \tag{30}$$

The "extinction" rate parameter θ'_B is assumed to be smaller than θ_B. In this the model follows the convention of similar linear models. Without such an assumption, such models could not generate conditioning under lean reinforcement schedules. In our simulations, θ'_B was roughly 1/6 of the value of θ_B, where B is any arbitrary event.

The ISI parameter $\tau(\Delta t)$ is the same as previously described. Here, Δt is the interval between A and non-B. In a sequence such as AABA, in which B had occurred at some earlier time, the first A predicts non-B simultaneously and $\Delta V_A^B < 0$. In the second bin, A also predicts non-B simultaneously, but this decrement might be discarded in favor of the one between the first A and the current non-B should this be a more favorable ISI. In the third bin, $\Delta V_A^B > 0$ by virtue of the most favorable ISI with respect to either the first A or the second (but not both). In the fourth bin, non-B is predicted by all three As, but only the one yielding the greatest decrement is retained.

The most important point to bear in mind regarding decrements in V_A^B due to non-B is that the salience of A is unaffected by these decrements until such time as A and B occur together in the same bin. At that time, the salience of B is also adjusted.

We do not assume that the nonoccurrence of an event has salience in its own right. This does not imply that organisms do not detect or notice that a stimulus has ceased or is absent, but merely that the nonoccurrence of an event is not regarded as a signal or predictor of other events in the model. Salience reflects associability, not noticeability, although the two are related. In deciding to preclude nonoccurrences of events as predictors, we are simply cutting the potential number of associative relationships in half with little apparent loss of predictive power.

A Single Stimulus A in Context X

The mechanics of the model may now be illustrated in detail. Each equation is stated for successive time bins. For convenience, the duration of each bin is equal to q, the optimal ISI for conditioning. Assume that the values of all associative relationships are initially zero.

Bin 1: Contains only X. All simulations assume that the first bin contains only X, the context. There are no computations on V_A^X, V_X^A, V_A^A, V_X^X, α_A, or α_X.

Bin 2: Contains X and A. Here, X has preceded A by one time bin. Since the bin equals the optimal ISI, $\tau = 1$. Hence, $\Delta V_X^A = \alpha_X \theta_A(\tau = 1)(\lambda_A - V_X^A)$. If X had not preceded A, ΔV_X^A would have been smaller by virtue of the less favorable simultaneous relationship between the two events. As indicated previously, the simulation chooses the maximum absolute value of ΔV_X^A and disregards the other.

Bin 2 does provide the initial increment in V_A^X, but the size of the increment is limited because $\tau = s$, where s is the small fraction associated with less than optimal ISIs. Thus,

$$\Delta V_A^X = \alpha_A \theta_X(\tau = s)(\lambda_X - V_A^X).$$

The relationship whereby A predicts itself (V_A^A) is unchanged in Bin 2 ($\Delta V_A^A = 0$) because an event cannot form an associative relationship with itself within a bin. However, the simulation permits that such a relationship can change from one bin to the next. Accordingly, $\Delta V_X^X = \alpha_X \theta_X(\tau = 1)(\lambda_X - V_X^X)$ because X occupies both the first and second bins. This equation holds even though X is continuous and uninterrupted. In essence, the model assumes that the longer an event is "on," the stronger the predictive relationship that implies its continuation.

Since X and A occupy Bin 2, the salience of each is altered. The net value of $\Delta\alpha_A$ is a weighted average of two components, $\Delta\alpha_A^X$ and $\Delta\alpha_A^A$. The component $\Delta\alpha_A^X$ decreases because $V_X^X > V_A^X$ at this point in time. That is, $\Delta\alpha_A^X = -\alpha_A C_A(D_A^X - D_X^X)/2\lambda_X$. Similarly, $\Delta\alpha_A^A = -\alpha_A C_A(D_A^A - D_X^A)/2\lambda_A$, even though V_A^A has yet to change from its initial value. By Equation 23,

$$\Delta\alpha_A = \frac{\phi_X \Delta\alpha_A^X + \phi_A \Delta\alpha_A^A}{\phi_X + \phi_A},$$

representing a net decrease of α_A.

The two components of $\Delta\alpha_X$ are positive because $V_X^A > V_A^A$ and $V_X^X > V_A^X$. That is,

$$\Delta\alpha_X^A = C_X(1 - \alpha_X)(D_A^A - D_X^A)/2\lambda_A$$

and

$$\Delta\alpha_X^X = C_X(1 - \alpha_X)(D_A^X - D_X^X)/2\lambda_X.$$

The net value of $\Delta\alpha_X$ is positive:

$$\Delta\alpha_X = \frac{\phi_X \Delta\alpha_X^X + \phi_A \Delta\alpha_X^A}{\phi_X + \phi_A}.$$

In a nutshell, a single occurrence of A in context X has enhanced the salience of X while reducing the salience of A. The mechanism of this result is that X preceded A, and therefore X was in position to become the better predictor of events in Bin 2.

Bin 3: Contains X but not A. Since A has occurred, but is not present in Bin 3, V_X^A is reduced by the expression

$$\Delta V_A^X = \alpha_A \theta_X'(\tau = 1)(0 - V_A^X).$$

The ISI parameter τ is equal to 1 because the τ is selected to maximize $|\Delta V|$; in this case, the relevant link is between X in Bin 2 and non-A in Bin 3.

Although A is not in Bin 3, V_A^X is increased by virtue of the forward conditioning arrangement between Bin 2 and Bin 3:

$$\Delta V_A^X = \alpha_A \theta_X(\tau = 1)(\lambda_X - V_A^X).$$

At this time, the other predictive relationship for A, that to itself, would be decreased by the following expression were it not for the fact that V_A^A is already zero.

$$\Delta V_A^A = \alpha_A \theta_A'(\tau = 1)(0 - V_A^A).$$

The predictive relationship between X and itself is increased in Bin 3 by the same equation for ΔV_X^X as in Bin 2. However, the magnitude of the increase is smaller because V_X^X is closer to asymptote (λ_X) than previously.

The saliences of A and X remain unchanged in Bin 3 because there are no computations on the various components of salience change, $\Delta\alpha_A^A$, $\Delta\alpha_A^X$, $\Delta\alpha_X^X$, and $\Delta\alpha_X^A$. As stated previously, changes in salience can occur only when two or more events occur simultaneously.

Bin 4: Contains X and A. The favorable ISI between Bin 3 and Bin 4 implies that V_X^A is increased by the amount

$$\Delta V_X^A = \alpha_X \theta_A(\tau = 1)(\lambda_A - V_X^A).$$

Less pronounced increases in V_X^A are also available from the more remote forward link from Bin 2 to Bin 4, for example, or from the simultaneous occurrence of X and A in Bin 4. As previously stated, the simulation discards these smaller increments in favor of the ΔV_X^A with the largest value of τ.

Unlike the situation that existed in Bin 2, where V_A^X increased by an amount consistent with simultaneous conditioning ($\tau = s$), the simulation may now choose a potentially larger value of τ bridging the interval from Bin 2 to Bin 4. This is some intermediate value ($\tau = i$), smaller than $\tau = 1$ but presumably larger than $\tau = s$. Accordingly, the maximum available increase in V_A^X is given by

$$\Delta V_A^X = \alpha_A \theta_X(\tau = i)(\lambda_X - V_A^X).$$

Similarly, V_A^A increases by the amount

$$\Delta V_A^A = \alpha_A \theta_A (\tau = i)(\lambda_A - V_A^A).$$

Provided α_A is initially greater than α_X, and provided the ISI is not too great, prediction of A would ultimately swing in A's favor with repetitions of the sequence of events described in this section, and $\Delta\alpha_A^A$ would become positive. As we shall see, increases in α_A due to the component $\Delta\alpha_A^A$ occur only when $\phi_A > \phi_X$ in Equation 23.

Finally, the two components of $\Delta\alpha_X$, $\Delta\alpha_X^A$, and $\Delta\alpha_X^X$, are both positive because at this point $V_X^A > V_A^A$ and $V_X^X > V_A^X$.

Latent Inhibition (LI)

The foregoing examples illustrate how the salience of an event such as a to-be-conditioned stimulus can decrease by repeated presentation in some context X. If event A is a potential CS, it is obvious that the model accounts for retarded conditioning (slow growth of V_A^{US}) in terms that are consistent with Mackintosh's (1975) theory. The bonus is that temporal parameters influence the process in a way that is made explicit in the present model. The effect of increasing the number of bins between successive presentations of A (henceforth referred to as the intertrial interval, ITI, to bring our terminology in line with customary usage by animal learning theorists—although ISI would be more appropriate) is to decrease α_A and enhance LI. This is because A is typically a weaker predictor of events than is X. With short ITIs, however, V_A^A can increase to a point at which α_A does not decline by repeated presentation in X. Given the proper initial value of parameters, the model predicts LI with long ITIs during the CS preexposure phase, but little or no LI with short ITIs.

Figure 1 (Panels A and C) shows the results of a computer simulation of LI in normal animals for both long and short ITIs. For each ITI, there is a CS-preexposure treatment and a control treatment such that the animal experiences X only during the preexposure stage. For these simulated animals, V_X^X increases to λ_X, but α_X remains at its initial value. The parameters were initialized as follows: $\alpha_A = .6$, $\alpha_X = .15$, $\alpha_{US} = 1.0$, $\phi_{US} = 1$, $\phi_A = .16$, $\phi_X = .01$, $C_{US} = C_A = C_X = .3$, $\theta_{US} = \theta_A = \theta_X = .06$, $\theta'_{US} = \theta'_A = \theta'_X = .01$, $\lambda_{US} = \lambda_A = \lambda_X = 1$. All associative relationships were initialized to zero.

The simulation shows a pronounced LI effect for the preexposed and control treatments at the long ITI, but virtually no LI at the short ITI. Solomon and Moore (unpublished observations) have preliminary results from rabbit NMR conditioning consistent with the simulation. The pattern also agrees with most of the available published reports (Lantz, 1973; Schnur & Lubow, 1976), which indicate LI only at ITIs over 30 sec and no LI at ITIs under 20 sec.

As a corollary of the model, increasing the duration of the CS, so that a single presentation extends over two or more bins of optimal duration should decrease the LI effect by the same mechanism as in the case of a short ITI. So far as we are aware, the appropriate experimental test of this prediction has yet to be carried out.

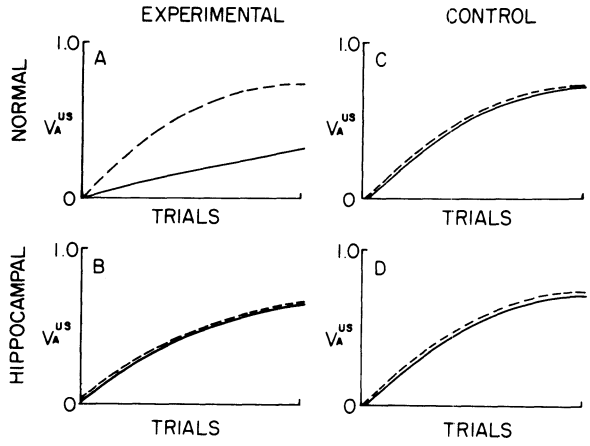

Figure 1. The changes in associative strength (V_A^{US}) are shown for the nominal CS (A) on a trial-by-trial basis, as a function of ITI (short or long), for experimental and control condition simulations of the latent inhibition (LI) procedure. The initial values for the following parameters of the model equations apply to all simulations: $\lambda_{US} = \lambda_A = \lambda_X = 1$; $\phi_{US} = 1$; $\phi_A = .16$; $\phi_X = .01$; $C_{US} = C_A = C_X = .3$; $\theta_{US} = \theta_A = \theta_X = .06$; $\theta'_{US} = \theta_A = \theta_X = .01$; $\tau(\Delta t < q) = 0$; $\tau(q) = 1$. The solid and dashed lines represent the values of association for A at long ITIs and short ITIs, respectively. The short ITIs are 2 time bins in length, and the long ITIs are 30 time bins in length. All graphs represent simulations of the acquisition stage of LI. The horizontal axes represent 50 trials of A-US pairings in a forward-delay arrangement. The initial salience parameters (α) for each simulation were derived from simulations of the preexposure condition (35 trials in length) appropriate to the group being simulated. Figure 1A simulates the normal experimental condition of the LI procedure, and the initial (α) and final (α') salience values for the two ITI conditions were: α_A(short) = .64; α_X(short) = .20; α'_A(short) = .96; α'_X(short) = .02; α_A(long) = .23; α_X(long) = .67; α'_A(long) = α'_X(long) = .42. Figure 1B simulates the hippocampectomized experimental condition, and the salience values were the following: α_A(short) = .71; α_X(short) = .26; α'_A(short) = .95; α'_X(short) = .26; α_A(long) = .67; α'_A(long) = .72; α'_X(long) = .91. Figure 1C simulates the normal control condition [α_A(short) = α_A(long) = .60; α_X(short) = α_X(long) = .15; α'_A(short) = .98; α'_A(long) = .96; α'_X(short) = .01; α'_X(long) = .02], and Figure 1D simulates the hippocampectomized control condition [α_A(short) = α_A(long) = .6; α_X(short) = α_X(long) = .15; α'_A(short) = α'_A(long) = .96; α'_X(short) = α'_X(long) = .18].

Role of Hippocampus in LI

The hippocampal formation is assumed to control the processes that decrease the salience of an event. Thus, if A is the target CS in an LI experiment, the effect of hippocampal disruption is to render all expressions of the form of Equation 6 equal to zero. Figure 2 (Panels B and D) shows the result of hippocampal disruption on LI at two ITIs. All parameters were initialized to the same values as in the simulation of normal LI in Figure 1. Figure 1 shows that

the only dramatic difference between normals and hippocampals is in the preexposed-long-ITI condition. Unlike normals, acquisition proceeds as though there had been no preexposures of the CS. With a long ITI, simulation nicely mimics the results of Solomon and Moore's (1975) demonstration that hippocampectomy disrupts LI.

We believe these simulations are instructive in illustrating not only how hippocampal disruption might be expressed in terms of a general effect on one class of transformations in the system of computations affecting associability, but also how the magnitude of the effect, and hence its detectability in a given situation, depends on the parameters of the situation. The simulations illustrated here show what might be expected when ITI is varied during CS preexposure. Other cases would illustrate that LI would occur even with a short ITI if either the salience of the context were initially greater than the salience of the CS or if the memorial representation (ϕ) of the context exceeds that of the CS. Hippocampal disruption would work against LI in these cases, but whether the effect would be noticeable would depend on the values assigned to parameters.

Blocking—Stage 1

A blocking experiment typically consists of a group of animals conditioned to one CS (A) and another group which experiences only the context in which the conditioning occurs. Both groups are then switched to a procedure in which A is compounded with another stimulus (B) and both are paired with the US. In order to illustrate our simulations of blocking, we must begin with a complete description of what happens during a single acquisition trial in the first stage of an experiment when the blocking group experiences the US and A together. During this stage, time is partitioned into bins containing subsets of three events, the context (X), the CS (A), and the US. As before, assume that the duration of a bin is equal to q, the optimal ISI in Equations 28 and 29. Acquisition trials were programmed to consist of two successive bins containing the CS, with the US occurring only in the second. This represents a forward-delay conditioning paradigm in which the CS and US terminate together. In order to set up the extinction operators and ϕ values, assume that the animal has already experienced all three events, X, A, and the US.

Bin 1: Contains X and A. We shall limit discussion to the fate of A. The associative value of A for predicting the US decreases in Bin 1 by the amount

$$\Delta V_A^{US} = \alpha_A \theta'_{US}(\tau = s)(0 - V_A^{US}).$$

However, V_X^A increases by the amount

$$\Delta V_A^X = \alpha_A \theta_X(\tau = s)(\lambda_X - V_A^X),$$

and V_A^A increases from previous experience with A at a time suitable for a modest increment given by

$$\Delta V_A^A = \alpha_A \theta_A(\tau = i)(\lambda_A - V_A^A).$$

The various components of $\Delta\alpha_A$ change as follows: Since the US does not occur in Bin 1, $\Delta\alpha_A^{US} = 0$. The decreases of V_A^{US} and V_X^{US} influence computation of $\Delta\alpha_A^{US}$ in the next bin. Assuming that $V_X^X > V_A^X$, $\Delta\alpha_A^X$ decreases (Equation 16), and assuming that $V_X^A > V_A^A$, $\Delta\alpha_A^A$ also decreases (Equation 20),

$$\Delta\alpha_A = \frac{\phi_X \Delta\alpha_A^X + \phi_A \Delta\alpha_A^A}{\phi_{US} + \phi_X + \phi_A} < 0.$$

Notice that ϕ_{US} contributes to the denominator of this expression even though the US does not appear in Bin 1. Assuming, as we have in our simulations, that the US is more strongly represented in memory than other events, the effect of adding ϕ_{US} to the denominator is to prevent α_A from being too greatly affected by momentary fluctuations among subordinate associations.

Bin 2: Contains X, A, and US. Once again limiting ourselves to the fate of α_A, V_A^{US} increases by the amount

$$\Delta V_A^{US} = \alpha_A \theta_{US}(\tau = 1)(\lambda_{US} - V_A^{US}).$$

The simulation disregards the slight contribution to V_A^{US} from simultaneous conditioning in Bin 2. A similar argument dictates that

$$\Delta V_A^X = \alpha_A \theta_X(\tau = 1)(\lambda_X - V_A^X)$$

and that

$$\Delta V_A^A = \alpha_A \theta_A(\tau = 1)(\lambda_A - V_A^A).$$

Despite these increments, V_A^X and V_A^A in all likelihood remain smaller than V_X^X and V_X^A, and therefore $\Delta\alpha_A^X$ and $\Delta\alpha_A^A$ are once again negative. However, the negative components are offset by the boost from $\Delta\alpha_A^{US}$. Here A is a better predictor of the US than is X because it has had fewer opportunities than X of being associated with nonoccurrences of the US. Moreover, since ϕ_{US} is typically much greater than ϕ_X or ϕ_A, $\Delta\alpha_A$ is positive by the expression

$$\Delta\alpha_A = \frac{\phi_{US} \Delta\alpha_A^{US} + \phi_X \Delta\alpha_A^X + \phi_A \Delta\alpha_A^A}{\phi_{US} + \phi_X + \phi_A}$$

Blocking—Stage 2

Stage 2 of a blocking experiment involves the introduction of another CS (B) to the two bins just described. Since V_A^{US} is higher than V_B^{US} at this point,

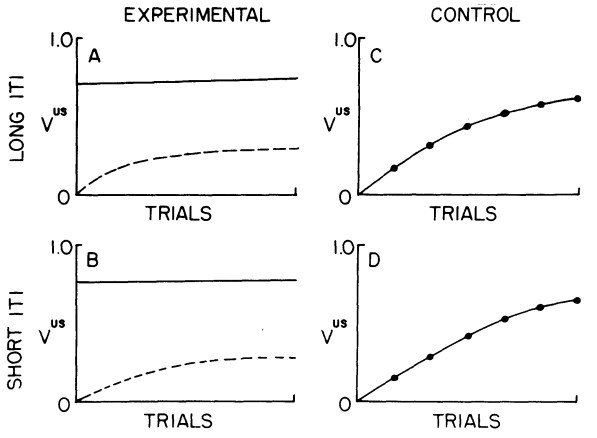

Figure 2. The changes in associative strength (V^{US}) are shown for the CSs (A and B) as a function of ITI for experimental and condition simulations of the blocking procedure. For the parameters that are constant across graphs (see Figure 1 legend), the initial values are the same as those of Figure 1. The initial values of the above parameters for the added CS (B) are the following: $\phi_B = .16$, $\lambda_B = 1.0$, $C_B = .30$, $\theta_B = .06$, $\theta'_B = .01$. All graphs depict a simulation of Stage II for a normal animal. The horizontal axes represent 50 trials of A+B/US pairings, and the ITIs are of the same duration as in Figure 1. Figure 2A simulates the experimental condition at the long ITI value ($\alpha_A = .98$, $\alpha_B = .60$, $\alpha_X = .007$, $\alpha'_A = 1.0$, $\alpha'_B = .02$, $\alpha'_X = .0009$), and Figure 2B simulates the experimental condition at the short ITI value ($\alpha_A = .98$, $\alpha_B = .60$, $\alpha_X = .01$, $\alpha'_A = 1.0$, $\alpha'_B = .01$, $\alpha'_X = .0003$). The solid lines represent the associative values for A and the dashed lines represent the values for B. Figure 2C simulates the control condition at the long ITI value ($\alpha_A = \alpha_B = .60$; $\alpha_X = .15$; $\alpha'_A = \alpha'_B = .58$; $\alpha'_X = .02$), and Figure 2D simulates the control condition at the short ITI value ($\alpha_A = \alpha_B = .60$; $\alpha_X = .15$; $\alpha'_A = \alpha'_B = .58$; $\alpha'_X = .01$). The beaded lines in Figures 2C and 2D represent associative values for both A and B.

it follows that α_B quickly falls (Equation 6) because all of the components of $\Delta\alpha_B$ are inferior predictive relationships to the other events. For example, the most heavily weighted component ($\Delta\alpha_B^{US}$) is negative because, although B is possibly in a more favorable predictive relationship to the US than X ($V_B^{US} > V_X^{US}$), it is in a less favorable relationship than is A ($V_A^{US} > V_B^{US}$). As indicated previously, the simulation chooses to adjust $\Delta\alpha_B$ according to the smallest discrepancies existing within the compound of stimuli. Thus, the fact that D_B^{US} may be smaller than D_X^{US} is of no consequence for the computation of $\Delta\alpha_B^{US}$ so long as $D_A^{US} < D_X^{US}$. As α_B decreases by the expression

$$\Delta\alpha_B = \frac{\phi_{US}\Delta\alpha_B^{US} + \phi_X\Delta\alpha_B^{X} + \phi_A\Delta\alpha_B^{A} + \phi_B\Delta\alpha_B^{B}}{\phi_{US} + \phi_X + \phi_A + \phi_B},$$

the opportunity for V_B^{US} to increase becomes stunted and blocking occurs.

Figure 2 illustrates a blocking experiment with normal animals at two levels of ITI. The simulation reveals that α_B quickly declines during Stage 2, and this loss of associability produces retarded growth of V_B^{US} (dashed lines) in comparison with the control groups shown on the right-hand side of the figure.

Role of Hippocampus in Blocking

Figure 3 summarizes our simulations of blocking in hippocampectomized animals. In agreement with Solomon (1979), there is no evidence of retarded growth of V_B^{US} in these animals as α_B does not decline because of the superior relationship of A over B in predicting the US. Figure 3 suggests that ITI does not play a major role in the unblocking effect.

Possible Implications for Goal-Directed Behaviors

The remaining question is how to formulate a response rule for cases in which more than one mode of responding or behavior is conditioned to the same reinforcer, each under the control of a different element of a stimulus compound. This is the problem confronted when trying to describe how attentional-associative models might be applied to complex learning tasks in which animals can adopt both cue and place strategies (cf. Black, Nadel, & O'Keefe, 1977; Moore, 1979b). A cue (A) and a place (B) can both be associated with reinforcement, but since they can be separated experimentally in any of a number of ways, place and cue strategies often lead to conflicting (mutually interfering) goal-directed behaviors. In general, A's relationship to the reinforcer might imply a strategy of approaching or avoiding A. Place B may carry the same implication to the animal, but if A and B are spatially separated, conflict and impaired performance is the result.

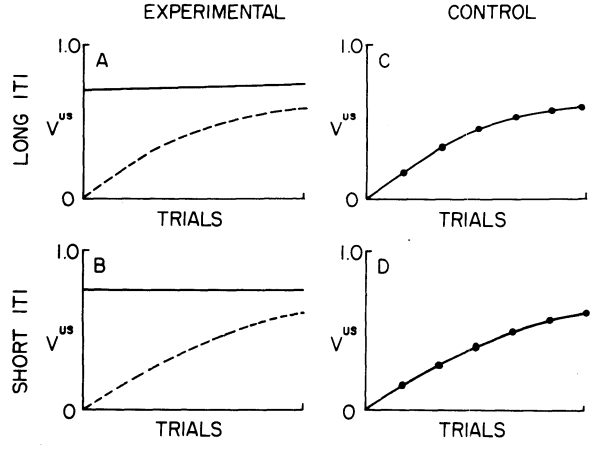

Figure 3. A blocking simulation identical to that of Figure 2 but with the assumption of hippocampectomy. The parameters that are constant across graphs (see Figure 1 legend) are identical to those of Figure 2. Control and experimental condition simulations are arranged as in Figure 2. The salience values for the graphs are as follows: Figure 3A, $\alpha_A = .98$, $\alpha_B = .60$, $\alpha_X = .18$, $\alpha'_A = 1.0$, $\alpha'_B = .60$, $\alpha'_X = .18$. Figure 3B, $\alpha_A = .96$, $\alpha_B = .60$, $\alpha_X = .18$, $\alpha'_A = .99$, $\alpha'_B = .60$, $\alpha'_X = .18$. Figure 3c, $\alpha_A = \alpha_B = .60$, $\alpha_X = .15$, $\alpha'_A = \alpha'_B = .61$, $\alpha'_X = .17$. Figure 3D, $\alpha_A = \alpha_B = .60$, $\alpha_X = .15$, $\alpha'_A = \alpha'_B = .60$, $\alpha'_X = .17$.

We believe that animals resolve such conflicts by adjusting the salience relationship between A and B in relation to the context X so that one associative value can gain a dominant position and thereby resolve the conflict. One way this could be done is by tuning out stimuli with comparatively large discrepancies with respect to the reinforcer by Equation 6 of the model. Animals with hippocampal damage, we suggest, cannot use Equation 6 or similar operations designed to reduce salience. The resulting distortion of their attentional-associative networks underlies their behavioral peculiarities.

REFERENCES

BLACK, A. H., NADEL, L., & O'KEEFE, J. Hippocampal function in avoidance learning and punishment. *Psychological Bulletin*, 1977, **84**, 1107-1129.

BLANCHARD, R. J., SHELTON, V. F., & BLANCHARD, D. C. Historical effects of stimulus exposure: Readiness to eat and object exploration. *Learning and Motivation*, 1970, **1**, 432-444.

GORMEZANO, I., & MOORE, J. W. Classical conditioning. In M. H. Marx (Ed.), *Learning: Processes*. Toronto: Macmillan, 1969.

GROVES, P. M., & THOMPSON, R. F. Habituation: A dual-process theory. *Psychological Review*, 1970, **77**, 419-450.

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

LUBOW, R. E., SCHNUR, P., & RIFKIN, B. Latent inhibition and conditioned attention theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 1976, **2**, 163-174.

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

MOORE, J. W. Brain processes and conditioning. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation: A memorial volume to Jerzy Konorski*. Hillsdale, N.J: Erlbaum, 1979. (a)

MOORE, J. W. Information processing in space-time by the hippocampus. *Physiological Psychology*, 1979, **7**, 224-232. (b)

O'KEEFE, J., & NADEL, L. *The hippocampus as a cognitive map*. Oxford: Oxford University Press, 1978.

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

SCHNUR, P., & LUBOW, R. E. Latent inhibition: The effects of ITI and CS intensity during preexposure. *Learning and Motivation*, 1976, **7**, 540-550.

SOLOMON, P. R. Temporal versus spatial information processing theories of hippocampal function. *Psychological Bulletin*, 1979, **86**, 1272-1279.

SOLOMON, P. R., & MOORE, J. W. Latent inhibition and stimulus generalization of the classically conditioned nictitating membrane response in rabbits (*Oryctolagus cuniculus*) following dorsal hippocampal ablation. *Journal of Comparative and Physiological Psychology*, 1975, **89**, 1192-1203.

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