

Conditioning and Learning Within Rule-Based Default Hierarchies

Keith J. Holyoak

University of California, Los Angeles

Kyunghee Koh

Richard E. Nisbett

University of Michigan

Abstract

We present a theory of classical conditioning based on a parallel, rule-based performance system integrated with mechanisms for inductive learning. Inferential heuristics are used to add new rules to the system in response to the relationship between the system's predictions and environmental input. A major heuristic is based on "unusualness": novel cues are favored as candidates to predict important, unexpected events. Rules have strength values that are revised on the basis of feedback. The performance system allows rules to operate in parallel, competing to control behavior and to obtain reward for successful prediction of important events. Sets of rules can form default hierarchies, in which exception rules "censor" useful but imperfect default rules, protecting them from loss of strength. The theory, implemented as a computer simulation, accounts for a variety of phenomena (e.g., rapid learning in certain inhibition paradigms, and failure to extinguish non-reinforced inhibitory cues) that previous associationist accounts have not dealt with successfully.

Introduction

Intelligence manifests itself in the adaptation of goal-directed systems to complex and potentially dangerous environments. The kinds of learning that underlie such adaptation fall under the rubric of *induction*, broadly defined as those inferential processes that expand knowledge in the face of uncertainty (Holland, Holyoak, Nisbett, & Thagard, 1986). Holland *et al.* presented a framework for induction that encompasses phenomena ranging from animal learning to human categorization, analogical reasoning, and scientific discovery. In the present paper we present a theory, derived from the general Holland *et al.* framework, that applies to some of the simplest forms of inductive learning: those observed in studies of classical conditioning in animals. The theory is implemented in a computer simulation that creates and revises rule-based default hierarchies. We will first describe some limitations of previous conditioning models.

CER Paradigm and Earlier Models

In a typical conditioning experiment, a rat is first trained to press a lever to get food. After many sessions of pressing the lever for food, a distinctive tone (the conditioned stimulus, or CS) is presented for several seconds. Just as the tone goes off, a shock (the unconditioned stimulus or US, also termed the "reinforcer") is delivered to the rat's feet. As this sequence of events is repeated, the rat soon begins to show signs of fear when the tone is heard. While the tone is on, the animal suppresses its routine lever pressing and eating, and displays a collection of behaviors, such as crouching, that constitute a *conditioned emotional response* (CER). The tone now signals shock and the rat exhibits fear in response to it.

Early theories of conditioning generally assumed that temporal contiguity of the CS and the US was necessary and sufficient to establish a conditioned response (CR). In fact, however, an event may be paired with another event and still not result in conditioning, as in the "blocking" phenomenon observed by Kamin (1968) (for recent reviews, see Mackintosh, 1983; Rescorla & Holland, 1982). Further difficulties for conditioning theories based simply on association arise in studies of *conditioned inhibition*. Kamin (1968) trained rats to associate noise with shock. He then paired the noise with light. This compound was presented for several trials, but was never paired with shock. Traditional associationist assumptions would predict that the light would take on the acquired fear conditioned to the noise. In fact, however, the effect of the light was to *inhibit* fear, which it did from the very first trial. The phenomena associated with conditioned inhibition indicate that whether conditioning is excitatory or inhibitory depends on the information that the cue provides about occurrence of the US, and cannot be predicted by simple temporal contiguity.

Rescorla and Wagner (1972) proposed an associationistic model of conditioning, in which the strength V_{CS} of an association is revised in accord with a linear model,

$$\Delta V_{CS} = \alpha_{CS}(\lambda_{US} - \sum_{j=1}^n V_j) \quad (1)$$

where α_{CS} is a constant that determines how fast conditioning can occur for a given CS, λ_{US} denotes the asymptotic limit of conditioning that can be supported by the US, and $\sum V_j$ represents the sum of the current strengths of associations to the US from the stimuli present (the particular CS plus all other concurrent cues).

Equation 1 is essentially equivalent to the Widrow-Hoff rule familiar in adaptive-systems theory, and is a generalization of the perceptron convergence rule (see Sutton & Barto, 1981). The Rescorla-Wagner model is thus closely related to a major class of strength-revision procedures used to model associative learning within adaptive networks of the sort currently being explored by connectionist theorists (Rumelhart, Hinton, & Williams, 1986). Within the connectionist framework, all learning is viewed as the product of incremental changes in connection strengths with experience.

Problems with the Rescorla-Wagner Model

Despite notable empirical successes, the Rescorla-Wagner model does not account well for a wide variety of other phenomena, some of which antedate their treatment and some of which are more recent. These inadequacies have spurred development of numerous other models that provide refinements and alternatives (Mackintosh, 1975; Pearce & Hall, 1980; Wagner, 1978, 1981). All of the models subsequent to that of Rescorla and Wagner (1972) adopt variations of the standard associationist framework in which conditioning is treated solely in terms of strength revision. In our view, none solves all of the empirical difficulties that beset the Rescorla-Wagner theory while preserving its successes. Here we describe some of these difficulties.

Learned irrelevance. A major problem for the Rescorla-Wagner formulation is that it does not account for changes in the processing of CSs with experience. Equation 1 predicts that a stimulus that is uncorrelated with a reinforcer will begin and end with zero associative strength. No distinction is drawn between stimuli that the animal has encountered before and those that it has not. But in fact conditioning is severely retarded if the to-be-conditioned stimulus has been presented previously (Baker & Mackintosh, 1977).

Rapid learning effects. None of the associationist models provides a satisfactory account of extremely rapid learning effects that are sometimes observed in conditioning studies. Kamin

(1968) and Rescorla (1972) have both shown that conditioning can take place after one trial, or indeed, even during an initial trial. For example, rats in the Kamin study showed less fear on the very first trial on which the excitatory stimulus was preceded by the inhibitory stimulus, despite the fact that the animal had no way of knowing whether this lessened fearfulness was justified. Such one-trial and no-trial effects look much more like inferences than like products of traditional trial-and-error learning.

Of course, learning is often much less rapid, and associationist models all include free parameters that govern rate of learning. The problem, however, is that the models provide no principled specification of when to expect rapid one-trial learning and when to expect slow learning requiring scores or even hundreds of trials. The rapid acquisition of conditioned inhibition is particularly problematic for all connectionist models of learning, because it involves the rapid alteration of a well-learned response (fear elicited by the excitatory CS), which has been gradually strengthened by reinforcement over many trials. There has been no demonstration that any connectionist model, using a consistent set of parameters for strength revision, can simultaneously account for the slow acquisition of a strong response with reinforcement, its typical pattern of slow extinction with nonreinforcement, and its immediate displacement by an incompatible response under certain specifiable conditions.

Conditioned inhibition. The Rescorla-Wagner model can account for the basic phenomenon of conditioned inhibition, as established in an A+, AX- paradigm (that is, stimulus A is always followed by the US, whereas the compound stimulus AX never is). The model assumes that strength values can be negative as well as positive. Cue A will reach an asymptotic strength equal to λ_{US} , whereas X will reach an asymptote equal to $-\lambda_{US}$; thus the net strength of the AX compound will be 0.

However, this account incorrectly predicts that inhibitory conditioning effects, once established, should be extinguishable by presenting the inhibitory CS alone, in the absence of either the US or of excitatory CSs. As Zimmer-Hart and Rescorla (1974) put it, "Assuming that nonreinforcement supports a zero asymptote . . . a simple nonreinforcement of a previously established inhibitor should produce a change. If V_X is negative, then the quantity $(0 - V_X)$ is positive and consequently V_X should be incremented toward zero when it is separately nonreinforced. That is, the theory predicts that repeated nonreinforced presentation of an inhibitor should attenuate that inhibition" (pp. 837-838). One might expect, however, that when a cue predicts the nonoccurrence of a US, then repeated presentations of the cue in the absence of the US will provide additional confirmations of the expectation of nonoccurrence, and hence would *enhance* its inhibitory properties. In fact (when ceiling effects on inhibition were controlled), this is the result obtained by Zimmer-Hart and Rescorla (1974).

Adaptation Within a Rule-Based Default Hierarchy

In view of the above and other limitations of associationist models of conditioning, it seems worthwhile to investigate an alternative approach. The theory of conditioning we will present is derived from the framework for induction proposed by Holland *et al.* (1986). We will first sketch the general framework, and then outline the specific theory of conditioning and describe its embodiment in a computer simulation.

Rule-Based Mental Models

Holland *et al.* proposed that representations of the environment take the form of sets of rules, with associated strength values, that comprise mental models. A mental model is an

internal representation that encodes the world into categories, and uses these categories to define an internal transition function that mimics the state changes that unfold in the world. In the relatively simple world of a rat's conditioning chamber, for example, the animal may learn that an occurrence of an instance of the category "loud tone" signals a transition to the environmental state "painful shock". The rat's knowledge about the relationship between tones and shocks might be informally represented by a rule such as, "If a tone sounds in the chamber, then a shock will occur, so stop other activities and crouch." Rules at different levels of generality form default hierarchies, as in situations that give rise to conditioned inhibition. For example, if tones are typically followed by shock, unless paired with a light, the rule "If tone, then expect shock" can serve as a useful but fallible default, to be overridden by the exception rule "If light and tone, then do not expect shock" when both rules are matched.

The Holland *et al.* (1986) framework provides a number of specific principles that can be applied to animal conditioning:

(1) The probability that a set of rules will control behavior increases monotonically with the strengths of the rules in the set relative to the strengths of competing rules that are matched. The function relating strength and response probability must be probabilistic in order to allow opportunities for weaker rules to be tested, and to gain strength if they prove more useful than their competitors.

(2) Rules form default hierarchies in which useful but imperfect default rules are "protected" from strength reduction by more specific exception rules that can override the defaults in particular circumstances.

(3) Only rules that succeed in controlling responses are subject to strength revision. In terms of an economic analogy, rules that control behavior "pay" for the privilege by a reduction in their strength, and must "earn" at least as much reward in the form of a subsequent strength increase in order to make the transaction worthwhile. Rules that do not gain control over responses produce no consequences for the system, and therefore neither gain nor lose strength.

(4) When multiple rules operate as a set to control behavior, they divide any attendant reward. This competition for reward implies that rules accrue greater reward when they uniquely make a correct prediction than when other rules make the same prediction (cf. Kelley's, 1973, "discounting principle"). Reward competition provides an inductive pressure that tends to favor general rules over redundant rules that are more specific, and that impairs learning of new rules that serve the same function as existing strong rules.

(5) New rules are generated in response to particular states of the system that suggest a new rule might be useful. In the current implementation of our conditioning theory, three triggering situations are identified: (a) the occurrence of an unexpected and important event; (b) the failure of a prediction based on a rule that had previously been highly successful; and (c) the occurrence of an unusual feature in temporal contiguity with a known predictor of an important event.

(6) Inferential heuristics favor certain features over others as building blocks for new rules. In particular, unusual features of the environment are favored as candidates to build the conditions of new rules.

A Rule-Based Theory of Conditioning

We have constructed a simulation model of conditioning based on the above principles. Figure 1 depicts the basic components of the processing system. In general terms, the system matches the conditions of rules against a "message" representing the current state of the environment (simulating perceptual input), and uses the matched rules to select an effector action and to generate a message describing the predicted next state of the environment. The predicted message is compared to that observed on the next time step, and the result governs the reward

given to rules that generated the prediction, storage of unusual events in a short-term buffer, and triggering of the generation of new rules.

Knowledge representation. As Figure 1 indicates, the model includes four types of information, which for simplicity we will describe in terms of four memory stores. Three of these store information of a declarative nature: a long-term store of all the cues that have appeared in the environment, tagged with a measure of degree of familiarity; a short-term store for recent unusual events; and currently active messages. The fourth store contains the rules in the system. Each of these stores is dynamically updated as the model operates in a simulated environment.

In the simulation, rules are represented in "classifier" notation (Holland, 1986). For purposes of exposition, we will use a more mnemonic notation. Thus the rule, "If tone occurs, then expect shock and crouch," will be represented simply as $T \Rightarrow S$. We will use the classifier symbol "#" to represent a maximally general condition. Thus the rule, "If in the conditioning chamber, expect shock and crouch," will be represented as $\# \Rightarrow S$. A rule with the condition "#" will be matched on every processing cycle.

The three declarative stores are each extremely simple. At any time, the message buffer contains three active messages: one describing the current environment, one describing the environment of the preceding cycle, and one, created on the preceding cycle, that predicts the current environment.

The short-term event buffer holds recent "unusual" events that occur in the environment. In the simulation, an unusual event is defined as an observed message that includes the onset of an "unfamiliar" feature, where familiarity of a feature is a function of its number of occurrences. This familiarity count for each feature is maintained in the long-term feature store. When an unusual event occurs, the message representing it, with the unfamiliar features tagged, is placed in the short-term event store and held for a few cycles, during which period it may be used by the rule-generation heuristics described below to form conditions of new rules.

Performance system. Each cycle of the processing system begins when a message describing the current state of the environment is received. The feature portion is compared with the corresponding portion of the prediction message posted on the previous cycle. The result determines how much reward, R (if any), is given to the rules that acted on the previous cycle. Specifically, reward is given in three circumstances: (a) if shock was predicted and occurred (a large positive reward); if absence of shock was predicted and shock did not occur (a lesser positive reward); and (c) if absence of shock was predicted but shock occurred (a negative reward, i.e., a punishment). Otherwise the reward is zero. The reward, if any, is added to the strength values of the relevant rules. The comparison is also used to trigger the generation of new rules (see below).

The message representing the current environment is compared to the message representing the previous environmental state to determine whether an unusual event has occurred. If a feature-onset occurs, a check of the long-term feature list is made to determine if the feature is unfamiliar, in which case the current event is defined as unusual and entered in the short-term event store.

Rule matching and response selection. Conditions of all rules are then matched against the message describing the current environment. Each matched rule posts a *bid*, which is a proportion of its strength. That is, the bid b made by Rule i will be

$$b_i = k * s_i, \quad (2)$$

where k is a constant between 0 and 1. Rules that make the same prediction sum their bids and act together as a set.

The rules that will govern the system's response are then selected in accord with Principles 1 and 2 above. Principle 1, the assumption of a probabilistic relationship between relative strengths of competing rules and response selection, is realized by applying a simple version of the Luce (1963) choice model to the strengths of matched rules. Principle 2, the assumption that default rules are protected by exception rules, is realized by allowing exception rules to "censor" their corresponding default rules (cf. Winston, 1986). When exception rules censor a default, the exception rules substitute for the default on that cycle. The rules that are selected to determine the system's response on a cycle will be termed the *winning set*, W . The action called for by the winning set is performed, creating a new predicted message, and the indicated effector action is taken. The next cycle then begins.

Strength revision. Strength revision takes place in two steps. In accord with Principle 3 above, only the rules in W have their strengths changed. The rules in W effectively compete for reward, as called for by Principle 4, in accord with the following scheme. First, when the winning set is selected, each rule in the set has its strength *reduced* by an equal portion of the summed bid made by the rules in W , $\sum_{j=1}^n b_j$, where n is the number of rules in W . Second, when the reward R is assigned to the winning set on the subsequent cycle, an equal portion of R (i.e., R/n) is added to the strength of each rule in W . The net change in the strength of Rule i , then, is

$$\Delta s_i = (R - \sum_{j=1}^n b_j)/n \text{ for } i \in W, 0 \text{ otherwise.} \quad (3)$$

Rule generation. The program contains three inferential heuristics for generating new rules, triggered by particular states of the system as specified by Principle 5. All three heuristics are specific instantiations of the unusualness heuristic (Principle 6).

(1) *Covariation Detection.* If a shock occurs unexpectedly, and is preceded by or concurrent with an unusual event stored in the short-term buffer, then a new rule will be constructed. The new rule will include the unfamiliar feature of the unusual event in its condition, and will have an action specifying expectation of a shock and crouching. For example, if an unfamiliar tone begins prior to an unexpected shock, the rule $T \Rightarrow S$ will be generated. If no unusual event is stored, and no other heuristic applies, then with some probability less than 1 a general rule is constructed with a maximally general condition and the same action as above. Thus if an unexpected shock occurs, and no unusual event is stored, the rule $\# \Rightarrow S$ may be generated.

(2) *Exception Formation.* If a strong rule makes an erroneous prediction about the presence or absence of a shock, and an unusual event occurred prior to or concurrent with the prior cycle (when the failed rule was matched), then exception rules are formed by (a) adding the unusual feature to the condition of the failed rule, and substituting the appropriate action, and (b) using the unusual features alone to form the condition. The failed rule is preserved as a default, tagged with the newly created exception rules. For example, suppose a tone occurs paired with an unfamiliar light, and the strong rule $T \Rightarrow S$ is a member of the winning set and creates an expectation of shock, which fails to occur. The rules $L + T \Rightarrow \bar{S}$ and $L \Rightarrow \bar{S}$ will then be generated. The former exception rule corresponds to the hypothesis that the unusual cue (L) signals nonoccurrence of shock only in the presence of the known predictor (T); the latter exception rule captures the more general possibility that the unusual cue might signal absence of shock regardless of whether the known predictor occurs.

(3) *Chaining.* If an unusual event is stored in the buffer, and a strong rule is included in the current winning set, then a new rule may be formed that uses the unfamiliar features of the

unusual event in the condition, and which has the same action as the parent rule. The initial strength is set to a proportion of the strength of the strong rule from which the new rule was constructed. Chaining tacitly seeks earlier predictors of the US; hence the strength of the new rule is set higher if the onset of the unusual event preceded the old CS (that is, preceded the cycle on which the parent rule was matched) than if it was concurrent with it. As an example, suppose a tone occurs concurrently with an unfamiliar light, and the strong rule $T \Rightarrow S$ is a member of the winning set. Then the rule $L \Rightarrow S$ may be generated (at the lesser initial strength value).

The three heuristics serve related but distinct functions. Covariation Detection provides initial rules to explain unexpected occurrences, whereas Exception Formation and Chaining build on existing partial knowledge. Exception Formation creates exception rules that may censor (and hence protect from further strength reduction) strong default rules that err under identifiable circumstances. Exception Formation is thus a reaction to an erroneous prediction. In contrast, Chaining represents an opportunistic attempt to identify an earlier predictor of a US that is already predicted by a CS.

Simulations of Conditioned Inhibition

We have simulated several variations of the CER paradigm, including blocking (Kamin, 1968) and the effects of statistical predictability on learning (Rescorla, 1972). The unusualness heuristic provides an explanation of the fact that unfamiliar cues are maximally conditionable. Here we will present simulations of two studies of conditioned inhibition described earlier, which pose difficulties for the Rescorla-Wagner model.

Rapid learning effects. An experiment by Kamin (1968) provides a dramatic demonstration of rapid inhibitory conditioning. Kamin trained rats for 16 trials to associate white noise with shock. He then created two different groups. Group LN received eight trials of a compound, simultaneous light-plus-noise stimulus that was never reinforced by shock, followed by four trials of the original noise stimulus which was again nonreinforced. Group N simply received 12 standard extinction trials, during which the noise was presented but never with shock. All animals received four trials per day.

Let us analyze the predictions our model makes for this study, based on rule generation and subsequent competition. Group N is of course expected to show just the customary gradual extinction as the rule $N \Rightarrow S$ dies a slow death due to nonreinforcement. The situation is much more complex for Group LN. Given that a trial in Kamin's experiment spanned a fairly long interval (several minutes), this time period would correspond to several cycles of matching and firing rules. Since no shock was presented on the first trial in which the light occurred along with the tone, the strong rule $N \Rightarrow S$ would repeatedly fail. Given the availability of an unusual event—the occurrence of the light—heuristics for rule generation will be triggered. In particular, Exception Formation should on the initial extinction trial generate new exception rules, $L + N \Rightarrow \bar{S}$ and $L \Rightarrow \bar{S}$. Both new rules will have an immediate inhibitory influence.

When the light is first paired with the excitatory noise, Chaining may create the excitatory rule $L \Rightarrow S$, which will compete with the other new rules. However, because the new cue does not occur prior to the original CS, the initial strength of the rule generated by Chaining will be low. Consequently, the influence of the new inhibitory rules will outweigh the influence of the excitatory ones, so that Group LN might be expected to show some inhibition of suppression even on the very first trial. Furthermore, the new inhibitory rules will of course be confirmed, and so the animal should show rapid development of inhibition over trials—much faster than Group N, which has no cue to suggest that the initial learning situation has now changed.

What should happen when, after the first eight trials, the noise alone is presented? For Group N, nothing interesting. This is merely a continuation of the slow competition between

the rule $N \Rightarrow S$ and its original competitor $\# \Rightarrow \text{Press}$. For Group LN, however, we expect a reversion to substantial suppression effects, because for rats in this condition the rule $N \Rightarrow S$ will have been protected to some extent due to the censoring effect of the successful exception rules.

Figure 2 presents the simulation results, which capture the major qualitative aspects of Kamin's findings. The data are presented as *suppression ratios*—the ratio of bar-pressing rate during CS presentations to the rate during the CS plus during its absence. A ratio of .5 indicates no excitatory conditioning to the CS, and a ratio of 0 indicates maximal conditioning. The results for Group N, presented with the noise alone, may be seen at the bottom of Figure 2. These animals showed the customary slow extinction process. The results for group LN are utterly different. The very first trial shows a substantially reduced suppression effect. The next trial, the first that confirms the new inhibitory rules, shows a further reduced suppression effect. By the fourth experience of the nonreinforced compound, the suppression ratio has become asymptotic.

Then, four trials after that, the single stimulus N is introduced. For Group N, this is by now simply the standard occurrence, but for Group LN it is an event not encountered since the original conditioning trials, during which N alone was always accompanied by shock. Because the rule $N \Rightarrow S$ has been partially protected from strength reduction by the exception rule $L + N \Rightarrow \bar{S}$, Group LN rats show considerable suppression on the very first presentation of N alone.

Increased inhibition due to an "extinction" procedure. We earlier discussed a phenomenon that is especially problematic for the Rescorla-Wagner formulation—failure to demonstrate extinction of the inhibitory power of an inhibitory cue that is presented in the absence of either the excitatory cue or reinforcement (Zimmer-Hart & Rescorla, 1974). In one experiment all rats were first given training in bar pressing to obtain food, followed by an initial session in which a 30-second tone was presented four times, ending each time with a shock. This would establish the rule $T \Rightarrow S$.

The animals then were divided into two groups, which received different procedures for inhibitory conditioning. For both groups, each subsequent session involved four presentations of the tone paired with shock, intermixed with four presentations of the tone in combination with a flashing light without shock. These events would generate the inhibitory rules $L + T \Rightarrow \bar{S}$ and $L \Rightarrow \bar{S}$ by Exception Formation, establishing the light as an inhibitory cue.

Group 1 received no other presentations of CSs. However, Group 2 also received four intermixed presentations of the light alone without shock. From the perspective of the Rescorla-Wagner theory, these were "extinction" trials that should have diminished the inhibitory power of the light, thus slowing down the acquisition of inhibition to the light-tone compound for Group 2 relative to Group 1. In contrast, from the point of view of our theory these are additional occasions for strengthening of the $L \Rightarrow \bar{S}$ rule. Since this rule contributes an inhibitory influence when the light-tone compound is presented, its strengthening when the light is presented alone should actually accelerate early acquisition of inhibition to the compound. However, the "light" rule will share reward with the other "light plus tone" exception rule when the compound is presented. Greater strength of the former rule will eventually lead to diminished strength of the latter. Accordingly, our model predicts that Group 2 will show less suppression than Group 1 to the compound early in acquisition, but that the two groups will show comparable suppression later in training.

Figure 3 presents the data from our simulation of this experiment. These data were obtained on test trials involving three reinforced presentations of the tone and two nonreinforced presentations of the light-tone compound. Suppression to the tone presented alone was asymptotic for both groups over the entire test period. The simulation model correctly indicates that animals in Group 2, which experienced separate presentations of the inhibitory light CS without reinforcement, should exhibit *increased* inhibition to the light-tone compound during early trials.

Conclusions and Future Directions

We are optimistic that the present theory can provide insights into forms of learning more complex than classical conditioning. The model can be extended to the acquisition of rule sequences by adding the "bucket brigade" algorithm for back-chaining strength to early rules in a sequence that eventually achieved a goal (Holland, 1986). We also hope that the theory will at some level prove relevant to understanding higher-level human cognition. Some recent work has begun to apply theoretical models derived from studies of animal learning to human categorization and decision making. Gluck and Bower (1986) found evidence that people weight cues to category membership in terms of their relative predictive power: the degree to which a cue is used to predict category membership is decreased by the presence of other more valid cues. This phenomenon can be accounted for by models such as that of Rescorla and Wagner (1972) and our own in which redundant cues compete to acquire strength. Interestingly, as Gluck and Bower point out, most current theories of human categorization do not display this property. Given that the present model exhibits both competitive learning and the capacity to represent nonindependent cues (an important aspect of human categorization), it may prove applicable to the analysis of performance in categorization tasks.

Our general aim has been to develop a model that integrates mechanisms of hypothesis formation and strength revision within a comprehensive performance system. Theories that invoke the notion of hypothesis generation often have been unduly restrictive, in our view, in assuming that hypotheses are entertained and tested serially, and ultimately rejected if any exceptions are found. By representing hypotheses as rules that can operate in parallel, take on continuous strength values, interact as defaults and exceptions, and actively compete to control behavior and to gain reward for predictive successes, a great deal of theoretical power is gained. The present theory makes extensive use of mechanisms for strength revision; however, the introduction of heuristics for rule generation is crucial in allowing us to account for phenomena, such as rapid acquisition of exception rules, that purely associationist mechanisms have not dealt with successfully. We suspect that the explanatory power of learning models based entirely on strength revision will prove to have limits, and that those limits will be found to lie short of a full account of classical conditioning in rats, far less of human cognition. A major limitation of current connectionist models of learning is that the proposed algorithms for adjusting connection weights become computationally intractable when the number of interconnected units grows large. Heuristics that propose plausible candidate rules can function to drastically reduce the effective size of the search space in which strength-revision procedures operate.

Acknowledgements

Preparation of this paper was supported by NSF Grants BNS-8615316 to K. Holyoak and SES-8507342 to R. Nisbett. Requests for reprints may be sent to Keith J. Holyoak, Department of Psychology, University of California, Los Angeles, California 90024.

References

- Baker, A. G., & Mackintosh, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning and Behavior*, 5, 315-319.
- Gluck, M. A., & Bower, G. H. (1986). Conditioning and categorization: Some common effects of informational variables in animal and human learning. In *Proceedings of the Cognitive Science Society Conference*.
- Holland, J. H. (1986). Escaping brittleness: The possibilities of general purpose machine learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*, Vol. 2. Los Altos, Calif. Kaufmann.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986) *Induction: Processes of inference, learning and discovery*. Cambridge, Mass.: Bradford Books/MIT Press.
- Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation*. Miami, Florida: University of Miami Press.
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, 28, 107-128.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. New York: Oxford University Press.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552.
- Rescorla, R. A. (1972). Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 6. New York: Academic Press.
- Rescorla, R. A., & Holland, P. C. (1982). Behavioral studies of associative learning in animals. *Annual Review of Psychology*, 33, 265-308.
- 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.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In Rumelhart, D. E., McClelland, J. L., & the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*, Vol. 1. Cambridge, Mass.: Bradford Books/MIT Press.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, 88, 135-170.
- Wagner, A. R. (1978). Expectancies and the priming to STM. In S. H. Hulse, H. Fowler, & W. K. Honig (Eds.), *Cognitive processes in animal behavior*. Hillsdale, NJ: Erlbaum.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms*. Hillsdale, NJ: Erlbaum.

Winston, P H. (1986). Learning by augmenting rules and accumulating sensors. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*, Vol. 2. Los Altos, Calif. Kaufmann.

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

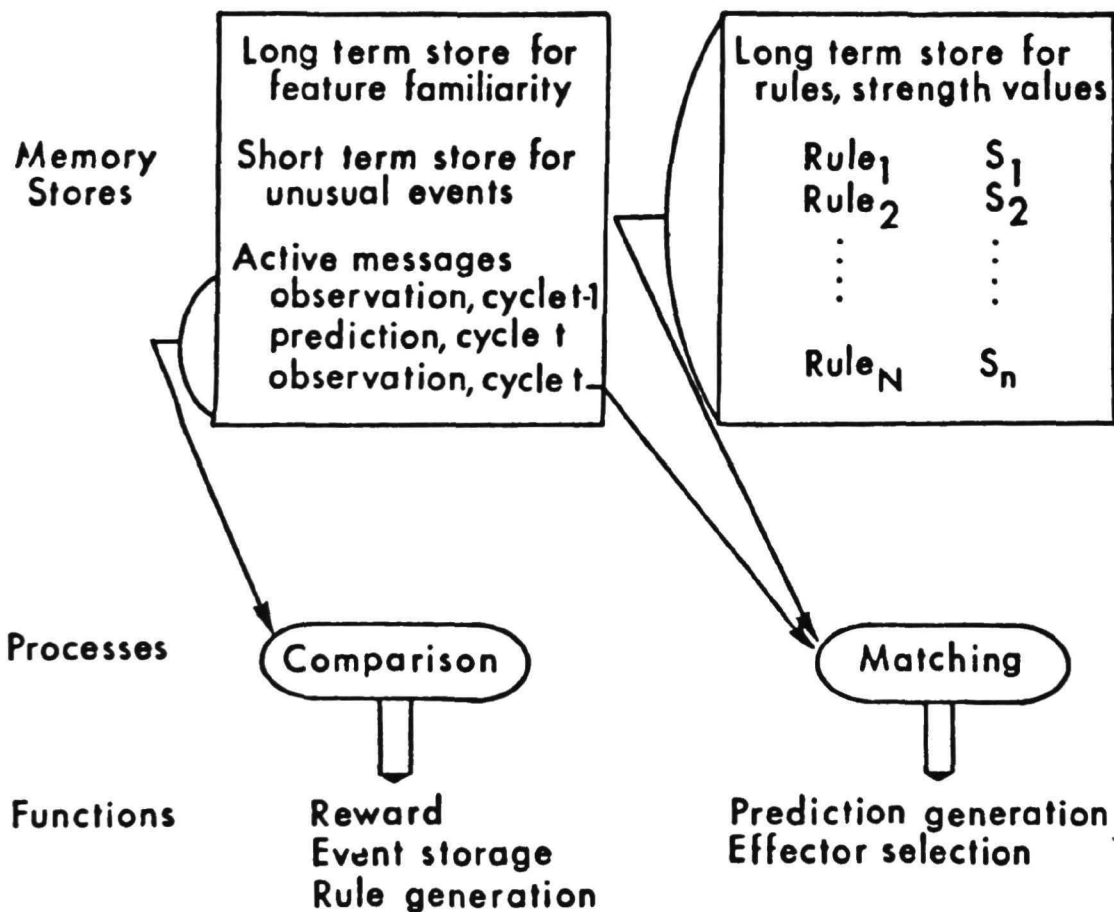


Figure 1. The memory stores and major processes involved in the process model of conditioning.

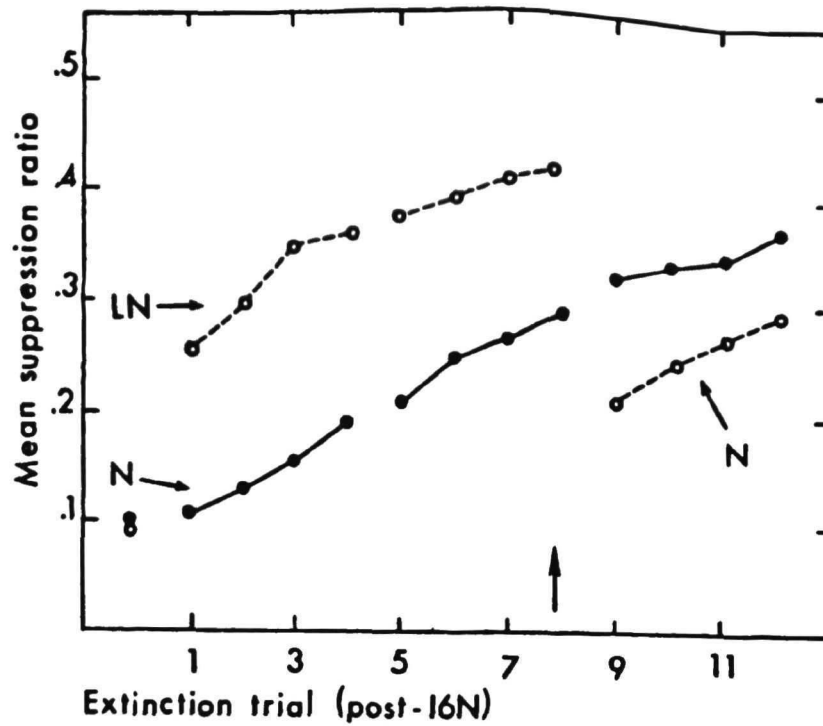


Figure 2. Simulation of experiment by Kamin (1968), showing predicted extinction of suppression, by trial, following 16 sessions of conditioning to noise (N). The groups were extinguished either to noise alone or to a light-plus-noise (LN) compound. The arrow in the abscissa indicates point at which group extinguished to compound was switched to noise alone.

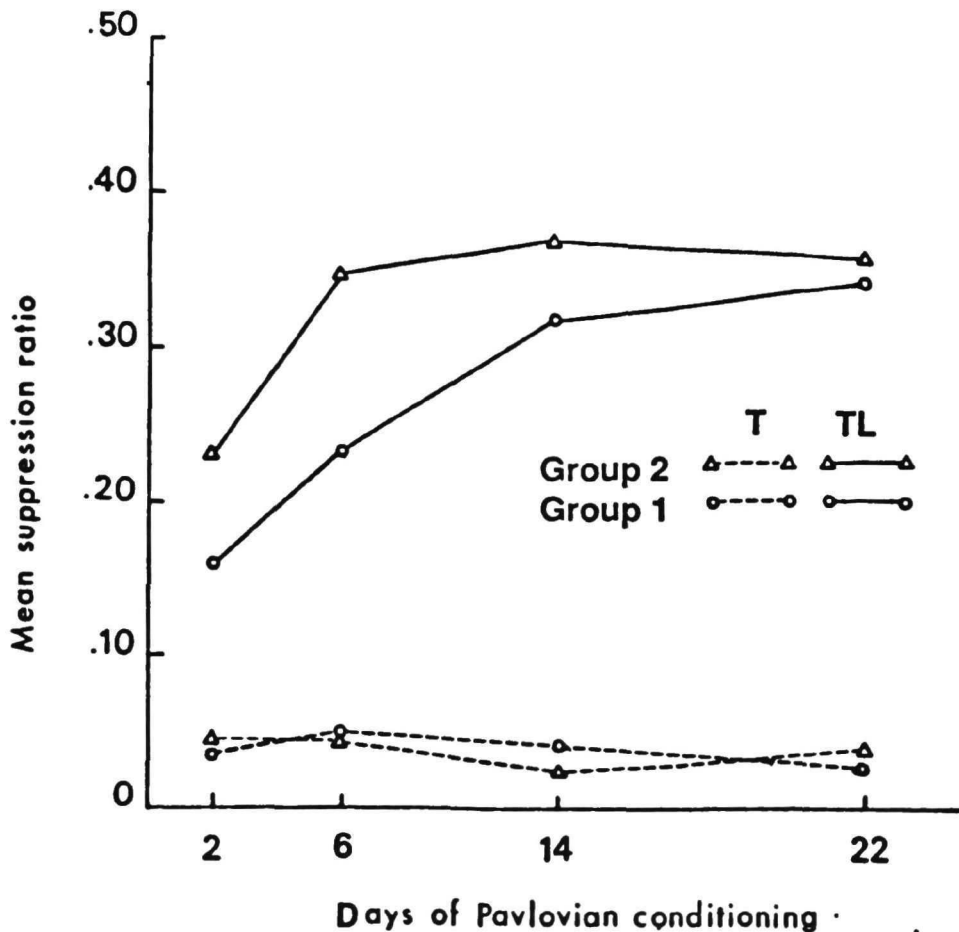


Figure 3. Simulation of experiment by Zimer-Hart and Reacorla (1974), depicting mean suppression ratios during tone-alone (T) and light-tone (LT) trials in single test sessions following various amounts of training. Both groups received T+, LT- presentations; Group 2 also received intermixed L- presentations.