SIMULTANEOUS CONFIGURAL CLASSICAL CONDITIONING

Jeffrey C. Schlimmer and Richard H. Granger, Jr.
Department of Information and Computer Science
University of California, Irvine

ABSTRACT

Humans and animals have the ability to learn complicated configurations of environmental cues that are predictive of important events. In classical conditioning, this task is called *configural* conditioning. Psychologists have studied this phenomenon since Pavlov's time, yet several of the contemporary learning models provide only partially satisfactory explanations. Most models provide mechanisms which select among possible predictive stimuli, but they fail to explicitly identify predictive *combinations* of stimuli and are thus restricted to learning only a relatively simple set of possible associations. In this paper we discuss a learning method which accounts for some configural conditioning results. Using an implemented system, we demonstrate the effectiveness of this method by modeling configural conditioning data from a pair of representative experimental studies.

INTRODUCTION

Consider the problem of trying to learn the precise configuration of weather cues that indicate rain. The appropriate description may include falling barometer readings and high humidity or a high degree of atmospheric ionization. The study of the ability of humans and animals to associate *sets* of stimuli with an important event is called *configural* conditioning. For example, in one case dogs were trained to associate a simultaneous presentation of six specific cues with delivery of meat powder (Razran, 1965). The animal then expectantly salivated only when all of the six cues were present and not when any subset of them were. Remarkably enough, a dog can also be trained to expect food only when one of two features occurs *separately* and not when they occur together (Woodbury, 1943). The ability to form associations with a number of Boolean combinations of features (conjunction, disjunction, exclusive-disjunction) is a well studied phenomenon in experiments on human ability (Grings, 1972; Bruner, Goodnow, & Austin, 1965) and in animal classical conditioning (Whitlow & Wagner, 1972; Saavedra, 1975).

Several contemporary animal learning theories, however, have difficulty explaining these results. The Rescorla–Wagner (1972) model, for instance, expresses the strength of a configural association as the sum of the associations for each cue. By adjusting the individual strengths, it *selects* the most predictive stimulus. A configuration consisting of any of a number of cues is easily accommodated. Conjunction and exclusive-disjunction, however, require that the individual cues have a qualitatively different associative strength than their combination. Within a sensory dimension, this difficulty is usually finessed by assuming that the co-occurrence of two stimuli (say a blue light and a red light) results in some new *resonant* property of the stimuli (say a purple light). Feature selection models like this one rely on resonant features in order to learn associations involving conjunctive and exclusive-disjunctive configurations. However, two stimuli may be configured from different modalities, as in the case where a tone and a light are reinforced alone but not together. It is unclear that any additional property of the stimuli is present (though perhaps one might wish to argue that some property of "two-ness" exists). Yet without it, traditional models cannot explain the effectual acquisition of these complex CSs; they assume that the association for a set of cues is simply the sum of associations with each part. A exclusive-disjunctive configuration requires weak learning of the compound to arise out of strong learning of each component. An additional difficulty of assuming the presence of resonant properties is that the number of these features must increase exponentially with the number of cues in the environment.

A number of artificial intelligence learning methods also have difficulties explaining these experimental data. A common assumption in concept attainment work, for instance, is that the identity of an instance can be determined via a conjunctive description of its features (Mitchell, 1982). A simple disjunctive description cannot be learned or represented in many cases, much less the exclusive-disjunction relationship. In this paper we present a model, STAGGER, which has the functional flavor of a feature selection model, but goes beyond this to form new, compound features. We further demonstrate its ability to correctly model the results from two representative experimental studies.

A LEARNING MODEL: STAGGER

The foundation of STAGGER is a distributed representation of association, composed of a set of dually weighted predictive features. During each trial, a cumulative expectation of the US is formed by utilizing the pair of weights associated with each feature. These weights are easily adjusted as learning progresses, and their mathematical interpretation mirrors basic results in learning. A secondary form of learning comes into play after expectation of the US fails; new features are introduced into the representation which are more general, more specific, or inverted Boolean functions of existing features.

At a higher level, associations in STAGGER are initially formed between primary, perceptual features in a similar manner to the processes in feature selection models. As conditioning progresses, compound features are formed internally which do not have an immediate correlate to individual perceptions. These new features are then part of the selection process, enabling STAGGER to learn complex configurations without relying on potentially nonexistent resonant features. The formation of effective feature compounds, then, is a central part of the process of configural conditioning in our model. In the following sections, we first describe STAGGER's associational representation and its feature selection processes. We then explicate its configural learning mechanisms.

Representation and Expectation

Associations are represented in STAGGER as a set of dually weighted features. The two weights for each feature capture positive and negative implication: one weight represents the sufficiency of the feature for the US, or $(CS \Rightarrow US)$, and the other represents its necessity, or $(\neg CS \Rightarrow \neg US)$. The mathematical measures chosen for these weights mirror the results of contingency experiments on learning in humans (Wasserman, Chatlosh, & Neunaber, 1983) and animals (Rescorla, 1968; Colwill & Rescorla, in press). Specifically, a novel cue comes to be excitatorily associated with an unpleasant stimulus only if the probability of the US in the presence of the CS is greater than its probability in the absence of the CS: $p(US|CS) > p(US|\neg CS)$. In behavioral terms, this means that if either the CS or the US frequently occurs alone, the subject still learns an association between the two stimuli. However, if they both occur alone even a few number of times, learning about their association is severely impaired.

With this in mind, STAGGER uses logical sufficiency (LS), or positive likelihood ratio, as a measure of sufficiency (Duda, Gaschnig, & Hart, 1979). Similarly, logical necessity (LN), or negative likelihood ratio, serves to measure necessity. They are defined as:

$$LS = \frac{p(CS|US)}{p(CS|\neg US)} \qquad\qquad LN = \frac{p(\neg CS|US)}{p(\neg CS|\neg US)}$$

LS ranges from zero to positive infinity and is interpreted in terms of odds. (Odds may be easily converted to probability $p = odds/(1 + odds)$.) An LS value less than unity indicates a negative correlation, unity indicates independence, and a value greater than unity indicates a positive relationship. LN also represents odds. However, an LN value near zero indicates a positive correlation, while a value greater than unity indicates negative correlation. For both LS and LN, unity indicates irrelevance. The LS and LN measures adhere to the contingency law, for it can be shown via algebraic manipulations that $LS > 1$ and $LN < 1$ if and only if $p(US|NC) > p(US|\neg NC)$ (Schlimmer, 1986).

In a given trial, all of the individual feature association weights influence US expectation. Following the mechanism used by Duda, Gaschnig, and Hart (1979), expectation of the US is the product of the prior odds of the US, the LS values of all present features, and the LN values of all absent ones.

$$Odds(\text{US}|\text{CSs}) = Odds(\text{US}) \times \prod_{\forall present} LS \times \prod_{\forall absent} LN$$

Table 1: Possible CS–US trial types.

	CS Present	CS Absent
US Present Reinforced	Confirming Positive (C_P)	Infirming Positive (I_P)
US Absent Nonreinforced	Infirming Negative (I_N)	Confirming Negative (C_N)

The resulting number represents the odds in favor of the US occurring. The effect of this multiplicative calculation is that learned associative strengths have a cumulative influence on US prediction. Two features which are highly predictive of the US cause a greater expectation when both are present than when only one of them is. However, as we will see in following sections, STAGGER is not confined to implicit representing configurations (via a summational effect of components) since it formulates new, explicit compound features which develop independent associative strengths.

Learning Mechanisms

In addition to computing a holistic expectation from the dual associational weights, STAGGER incrementally modifies the feature weights and generates new features. These two latter abilities allow STAGGER to adapt its associational description to better reflect the conditioning environment.

Feature selection

The sufficiency and necessity weights associated with each of the features may be easily adjusted. Consider the possible situations that arise during a conditioning trial. Following the terminology used by Bruner, Goodnow, and Austin (1956), a reinforced trial is positive evidence which may either *confirm* the predictiveness of a feature (if it is present in this trial) or *infirm* the feature's predictiveness (if it is absent). Similarly, a nonreinforced trial is negative evidence which either confirms an absent feature or infirms a present one. Table 1 summarizes these possibilities. In terms of these matching events, the contingency law implies that learning occurs in cases involving at most one type of infirming evidence. In situations with even small amounts of both positive and negative infirming evidence, subjects fail to learn an association.

LS and *LN* may be calculated by keeping counts for each feature of the possible situations listed in Table 1.

$$LS = \frac{C_P(I_N + C_N)}{I_N(C_P + I_P)} \qquad\qquad LN = \frac{I_P(I_N + C_N)}{C_N(C_P + I_P)}$$

The prior odds for the US are estimated as $(C_P + I_P)/(I_N + C_N)$. Note that the *LN* measure will rank features with negative infirming evidence highly, but features with both

Table 2: Feature formation heuristic.

Expectation	Reinforced	Error type	Feature formed
US	¬US	Commission	`AND[f1,f2]`
¬US	US	Omission	`OR[f1,f2]`
—	—	Either	`NOT[f1]`

types of infirming evidence poorly. This reflects learning in partial reinforcement situations (Fitzgerald, 1963) and is consistent with contingency experiments (Rescorla, 1968).

If STAGGER limited its learning to adjustment of the feature weights, the distributed association would be sufficient to accurately describe the class of "linearly separable" concepts (Hampson & Kibler, 1983). In this respect STAGGER is similar to *connectionist* models of learning when those models do not have any "hidden" units. The purpose of the hidden, internal units is to allow the encoding of more complicated associations. Feature formation processes in STAGGER serve an analogous purpose: individual features are combined into more complex Boolean functions.

Feature formation

STAGGER is not limited to acquiring summational associations between immediate, perceptual features. New, internal compound features are introduced by the model, allowing it to encode the potentially complex associations involved in configural conditioning. Because it is able to identify effective featural combinations internally, no assumption regarding additional resonant features is required. STAGGER follows three levels of heuristics in its formation of compound, internal features, and it constructs them using conjunction, disjunction and negation.

The first heuristic suggests a new compound feature when STAGGER makes an expectation error: either expecting the US in a nonreinforced trial or failing to predict the US in a reinforced trial. In the first case, the commission has admitted one too many possible situations as predictive of the US. A compound feature with a restricted application is formed using conjunction. This feature will be true less often than its components and can act to dampen the expectation process. In the second case, STAGGER is failing to include stimuli which do lead to the US, or making an error of omission. A more admitting, or more general, feature compound is formed using disjunction; it will be true more often than its components and thus loosens the class of possible predictors of the US. In either case, STAGGER forms a negated feature compound. Table 2 summarizes this heuristic.

Choosing appropriate features for new formations is accomplished via two additional heuristics. One heuristic *nominates* either present or absent features for combination, and the other narrows the possible features down by *electing* one or two of the most predictive. STAGGER's nomination heuristic specifies whether present or absent features are to be used in forming compound features, depending on the type of feature combination and prediction error. After STAGGER has made an error of commission, features present on

Table 3: Nomination heuristic.

Error type	Function formation	Feature nomination
Commission	AND[f1,f2]	Present, Absent
	OR[f1,f2]	Absent, Absent
	NOT[f1]	Present
Omission	AND[f1,f2]	Present, Present
	OR[f1,f2]	Present, Absent
	NOT[f1]	Absent

this nonreinforced trial may be partially necessary, but are clearly not sufficient for reinforcement. Conjunction nominates two necessary features, and thus a present feature is combined with an absent one. Nominating a present feature is motivated by noticing that some feature was present and suggested that this trial was likely to be reinforced. Disjunction nominates two sufficient features, so two features absent in this nonreinforced trial are chosen; no sufficient features were present. Negation is used to identify safety-signal features (those which, when present, indicate safety from the US) and thus nominates its component from the collection of features which were present. The appropriate nominations following an error of omission (an unpredicted, reinforced trial) are derived by similar reasoning. Table 3 summarizes the nomination heuristic.

The election heuristic further narrows the possible features for combination. Consider a situation leading STAGGER to appropriately form a new conjunction. For example, the familiar concept *father*: a parent and a male. The two features (parent and male) are always reinforced (father) though each is separately nonreinforced (a brother is male). This is negative infirming evidence (Table 1), and therefore LN which tolerates negative infirming evidence is used to elect features for a conjunctive configuration. By a similar argument LS elects features to be used in forming a disjunction. Features are elected equally by both measures for negated formations. Table 4 summarizes this third heuristic.

Table 4: Election heuristic.

Function	Election measure
AND[f1,f2]	$LN(\mathtt{f}i) \ll 1$
OR[f1,f2]	$LS(\mathtt{f}i) \gg 1$
NOT[f]	$LN(\mathtt{f}) \gg 1$ or $LS(\mathtt{f}) \ll 1$

These three heuristics may be used in concert, each one driving the others. For example, in the case of an unexpected, reinforced trial (an error of omission), a disjunction may be formed (see Table 2) to combine the present feature (Table 3) which has the lowest LN value (Table 4) with the absent feature with the lowest LN value. In many cases, these three heuristics cooperate in just this manner. However, there are some situations (as in

Table 5: Configural training.

HL+, H−, L−	Positive patterning or	AND[H,L]
HL−, H+, L+	Negative patterning or	XOR[H,L]

Woodbury's (1943) negative patterning discussed below) in which the election heuristic cannot offer any guidance. The process of forming effective internal, compound features then proceeds using the remaining two heuristics.

The feature formation process is limited by pruning ineffective compound features. Specifically, LN is used to assess the validity of a new conjunction. This more restrictive feature is true less often that its components are (it is guaranteed to have the same or less negative infirming evidence), so if it also has less positive infirming evidence (and has a better LN weight), it is deemed an effective feature. LS is used to test a more inclusive compound feature. Inverted features are tested by comparing them to the inverse of their Bayesian measure (e.g., $1/LS$). The role of these pruning measures is similar to that of the test component of a generate-and-test algorithm. The three formation heuristics serve to guide the generation of new feature compounds while the Bayesian measures are used to prune ineffective ones.

STIMULUS PATTERNS

Woodbury (1943) was one of the first American researchers to investigate learning of configural cues. In classical conditioning experiments with dogs, he studied different configurations of a low and a high buzzing sound which served as CSs for the delivery of a food pellet. He investigated two simultaneous configurations: positive patterning (in which only the presence of both buzzing sounds was reinforced), and negative patterning (where only the presence of either of the buzzing sounds was reinforced). In the first situation, a Boolean conjunction of the two buzzing sounds was reinforced; in the second, a exclusive-disjunction was reinforced. Table 5 summarizes these two training conditions.

Figure 1 depicts STAGGER's acquisition of the positive pattern (conjunctive) configuration. The upper heavy line represents conditioned responses (CRs) in trials which contained both of the buzzing sounds. The lower, light, solid line represents CRs to trials which contained only the lower of the two buzzing sounds and the lower, dotted line represents responses to the higher of the two buzzers. Each point represents the percentage of CRs in the last ten trials and is an average over ten separate program executions.

A strict feature selection model could learn to correctly predict reinforcement if there were a resonant feature resulting always and only from the co-occurrence of the low and high buzzers (Rescorla & Wagner, 1972, p. 86, fn. 2). Such a resonant feature is plausible given that both of the stimuli are within the same sensory modality. However, for the purpose of demonstrating the capabilities of the feature formation processes in STAGGER,

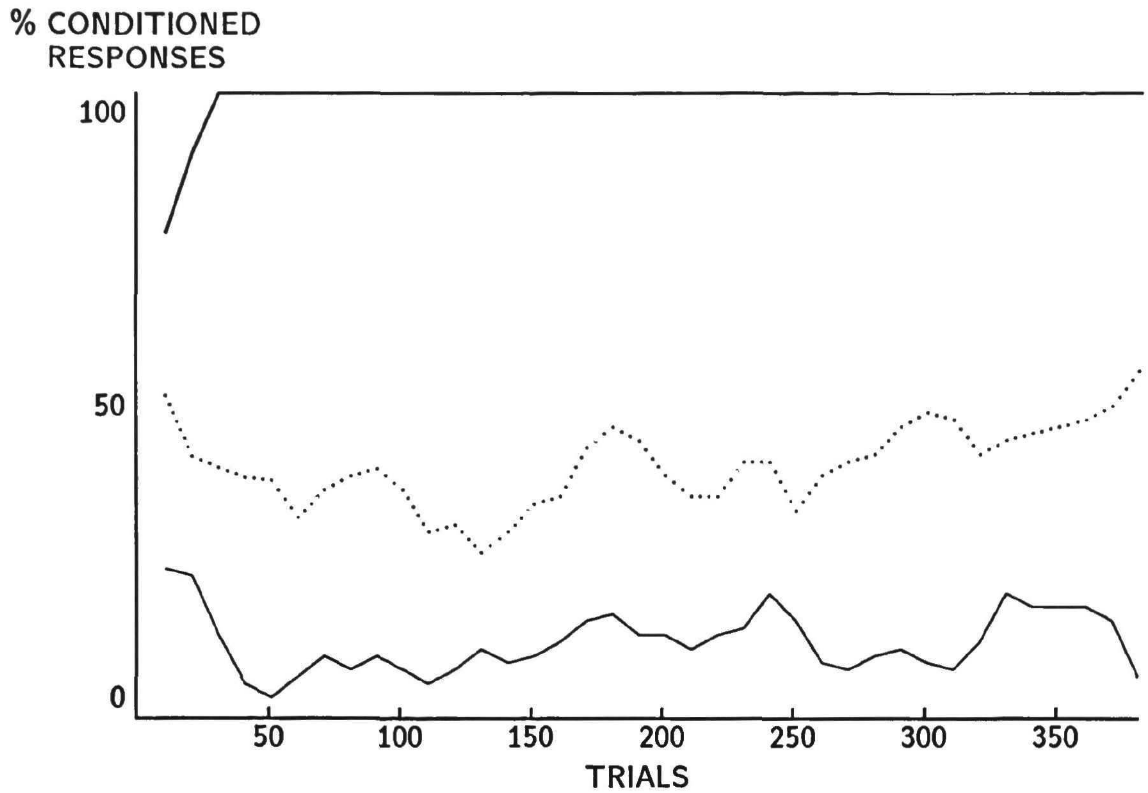

Figure 1: Conditioning to HL+, H−, L−.

this assumption about resonant features was omitted from the trial specification input to the program. Distinction between the combination of the two buzzers and either separately is facilitated by the introduction of the compound feature `AND[low-buzz,high-buzz]` via the heuristics described in section : the *LN* measure ranked low-buzz and high-buzz as the most effective individual predictors in an expected, but nonreinforced trial. A conjunctive feature was constructed using the low-buzz feature which was present and the high-buzz feature which was absent.

Figure 2 depicts STAGGER's acquisition of the negative pattern (XOR) configuration. Again the heavy line represents CRs to a co-occurrence of the two buzzing sounds; the light line, CRs to trials with only the low buzzing sound; the dotted line, CRs to the higher buzzing sound. After approximately 200 trials, STAGGER is effectively distinguishing between reinforced and nonreinforced trials.

If the presence of a resonant buzzing feature was assumed, a strict feature selection model could model this learning, for the resonant low and high buzzing feature could have a strong negative associative strength which would overpower either of the positive strengths of each of the individual buzzing cues.

STAGGER formed two compound features in order to accurately predict reinforcement: `AND[low-buzz,NOT[high-buzz]]` and `AND[NOT[low-buzz],high-buzz]`. A simple disjunction of these compounds captures negative patterning. However, the formation of these configurations was not as straightforward as was the case in Figure 1. As we inti-

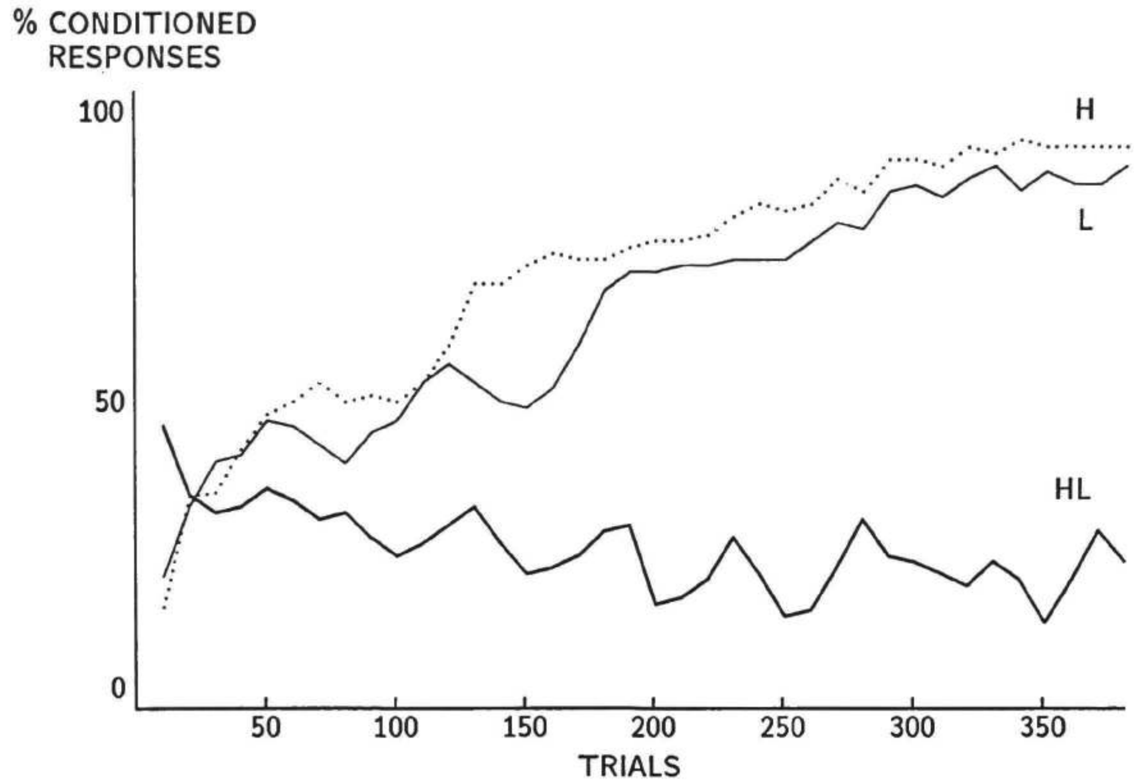

Figure 2: Conditioning to HL−, H+, L+.

mated previously, the *LS* and *LN* measures were unable to elect any effective features, for there was an equal amount of positive and negative infirming evidence for each of the cues. Therefore, exploration for predictive features occurred without Bayesian guidance. On an unexpectedly reinforced trial, for instance, STAGGER added three new compound features: a new disjunctive feature composed of a randomly elected, present feature and a randomly elected, absent feature; a new conjunctive feature made of a pair of randomly elected, present features; and a negation of a randomly elected, absent feature. This potential explosion of exploratory features was still subject to pruning via the Bayesian measures, though, and if each did not outperform the components from which it was composed, it was pruned.

BICONDITIONAL DISCRIMINATION

Saavedra (1975) has also studied simultaneous configural conditioning. However, unlike Woodbury, in each configuration all stimuli were from different sensory modalities. Assuming the presence of resonant features arising from the co-occurrence of features is therefore less reasonable. Instead of only two features, she utilized four, in pairwise configurations such that each feature was present in reinforced as well as unreinforced trials. No property of the features such as "two-ness" would aid in predicting reinforcement. One experimental group was given reinforced presentations of a tone (auditory cue A_1) and flickering light (L_1)

Table 6: Biconditional and component discrimination training.

A_1L_1 +, A_2L_2 +, A_1L_2 −, A_2L_1 −	Biconditional
A_1L_1 +, A_1L_2 +, A_2L_2 −, A_2L_1 −	Component

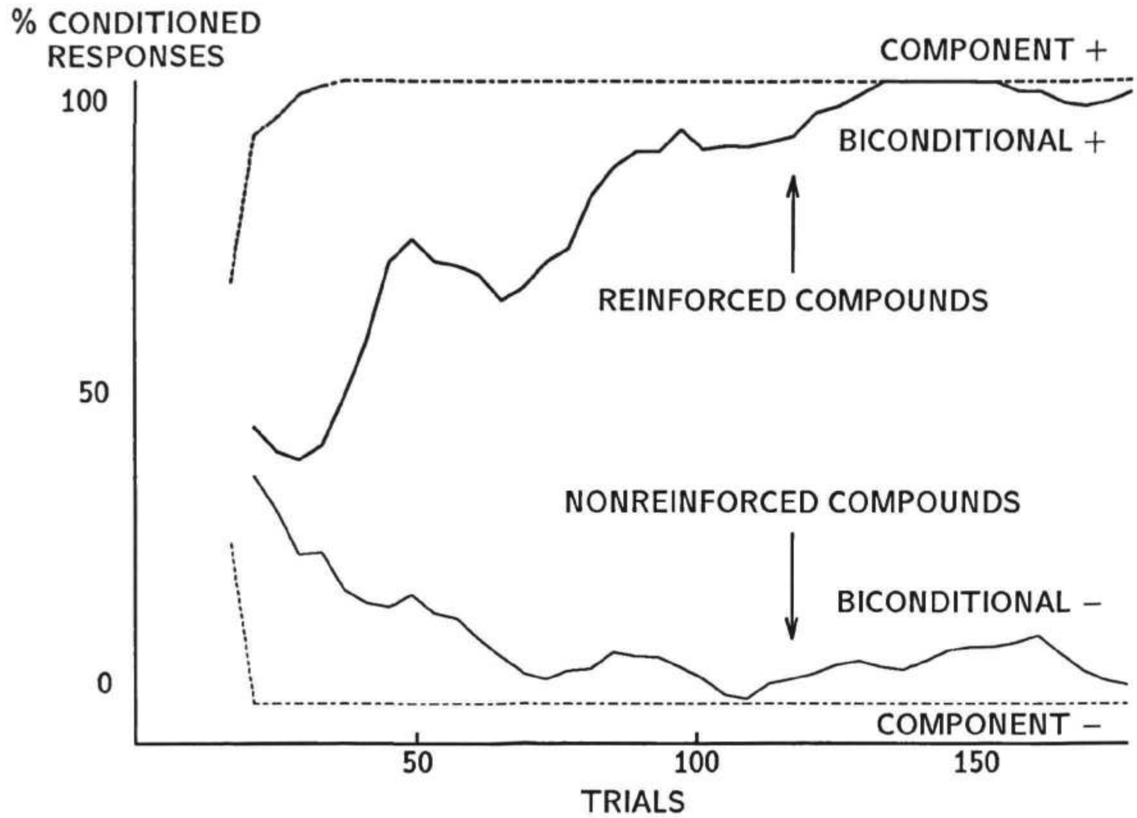

Figure 3: Biconditional versus component discrimination.

or a clicker (A_2) and a steady light (L_2). The alternate combinations were nonreinforced. This training is termed *biconditional discrimination* since reinforcement is conditional on two cues. For comparison, she also tested a simple component discrimination case where the tone was always reinforced. These training schedules are summarized in Table 6.

This experimental manipulation taxes the feasibility of a strict feature selection model since it is unlikely that the necessary resonant features are available; this class of models would predict that such an association would be unlearnable. Animal subjects, however, do learn the biconditional discrimination. Figure 3 overlays STAGGER's performance on both the biconditional and component discrimination cases. Each line represents the average percentage of CRs over ten separate program executions. The upper lines represent responding to the reinforced configurations; the lower lines, the unreinforced configurations. The solid lines represent conditioned responding in the biconditional discrimination case; the dashed lines correspond to component discrimination training.

The component discrimination training proceeds much more rapidly than the biconditional discrimination because the appropriate stimuli need only be selected in the former case, rather than formed, as in the latter case. In the biconditional discrimination task, STAGGER first forms the compound features `AND[A₁,L₁]` and `AND[A₂,L₂]` which are then used in the selection process. Without resonant features arising from the co-occurrence of A_1 and L_1, and A_2 and L_2, a strict feature selection model would be unable to acquire the biconditional discrimination. Associative strengths would have to be high enough for A_1 and L_1 to sum for a positive prediction when they occurred together, but low enough so that when A_1 and L_2 co-occurred, nonreinforcement would be expected. This is clearly impossible.

DISCUSSION

The two representative configural conditioning experiments of Woodbury (1943) and Saavedra (1975) indicate that animals are able to form associations between complex CSs and a US. Two categories of models have been proposed to account for this type of learning: feature selection only, and feature selection plus feature formation.

Feature selection models assume that the association accrued to a stimuli are summed when they co-occur (Rescorla & Wagner, 1972). The associative strength of a configuration of stimuli is simply the sum of the associative strengths of its components. A secondary assumption is that when two stimuli (say A and B) are present that a third *resonant* stimulus is present which has some of the properties of both (represented by AB); this AB stimulus is present always and only when both A and B are. A exclusive-disjunctive configuration like Woodbury's negative patterning is represented by a strong negative associative strength for the resonant feature and weaker positive associations for each of the components. This assumption extends the representational ability of feature selection models to include all possible Boolean functions.

There are two unsatisfactory consequences of assuming the presence of resonant features. First, while it seems plausible to assume that the simultaneous presence of a red light and a blue light adds a feature not present when either are presented separately (a purple light), this assumption seems tenuous when the stimuli are from different sensory modalities. Secondly, the number of stimuli from which the model must select grows exponentially with the number of stimuli that may be configurally associated. For example, if there are three stimuli (A, B, C), there must be four supplementary stimuli (AB, AC, BC, ABC) in order to select any configuration. Razran (1965) reports on experiments where six simultaneous features were conjunctively configured; 57 additional resonant features would be required if a feature selection model were applied. In general the number of resonant stimuli required by these models is $2^n - 1 - n$, where n is the number of perceptual, stimuli. Requiring the model to choose between $O(2^n)$ stimuli may be computationally infeasible.

The alternative we present here is a secondary process which formulates plausibly predictive compound features as they are needed. The number to be examined is therefore limited to those necessary and there is a corresponding reduction in computational load on the feature selection process. The fact that both approaches rely on the ability to form an association to a combination of cues is not new. However, unlike the strict feature selection model, STAGGER employs a feature formation component which can be used to configure individually perceptible cues into an explicit compound feature usable for learning. Instead of assuming that this process is already performed by the perception system via resonance (and its entailing assumptions of cross-modality resonance and exponential requirements), we prefer the property of necessity-driven feature formation. In this we concur with Razran when he notes:

> What seems more warranted is the view that, inasmuch as configures are formed and deformed through learning, their role is much more a function of the organism's conditioned past than of its sensory present, and, moreover, that their learning reveals the dynamic essence of their "becoming" if not also of their being (Razran, 1965, p. 244, fn. 3).

FUTURE WORK

One phenomenon unexplained by previous work on feature selection is that of sequential configural conditioning. In these types of experiments, effectively predicting the US requires discerning a sequential configuration of the cues in the environment. Woodbury (1943) also trained dogs to expect a food pellet only when the low buzzing sound followed the high buzzing sound. While we have provided some explanation of mechanisms which could give rise to simultaneous configural conditioning, we have yet to address the larger issue of associating sequences with outcomes. We believe that a featural formation approach, where sequences are constructed and their effectiveness evaluated, will prove useful.

ACKNOWLEDGEMENTS

This research was supported in part by the Office of Naval Research under grants N00014-84-K-0391 and N00014-85-K-0854, the National Science Foundation under grants IST-81-20685 and IST-85-12419, the Army Research Institute under grant MDA903-85-C-0324, and by the Naval Ocean Systems Center under contract N66001-83-C-0255. We would like to thank Michal Young for his early involvement with this project and the machine learning group at U.C. Irvine, who collectively have helped clarify a number of interesting issues and have provided a supportive environment for the exploration of new ideas.

REFERENCES

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1965). *A study of thinking.* New York: John Wiley & Sons, Inc.

Colwill, R. M., & Rescorla, R. A. (in press). Associative structures in instrumental learning. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory.* New York: Academic Press.

Duda, R., Gaschnig, J., & Hart, P. (1979). Model design in the Prospector consultant system for mineral exploration. In D. Michie (Ed.), *Expert systems in the micro electronic age*. Edinburgh: Edinburgh University Press.

Fitzgerald, R. D. (1963). Effects of partial reinforcement with acid on the classically conditioned salivary response in dogs. *Journal of Comparative and Physiological Psychology, 56*, 1056–1060.

Granger, R. H., Jr., & Schlimmer, J. C. (1985b). Learning salience among features through contingency in the CEL framework. *Proceedings of the Seventh Annual Conference of the Cognitive Science Society* (pp. 65–79). Irvine, California: Lawrence Erlbaum Associates.

Grings, W. W. (1972). Compound stimulus transfer in human classical conditioning. In A. Black & W. F. Prokasy (Eds.), *Classical conditioning II*. New York: Appleton-Century-Crofts.

Hampson, S., & Kibler, D. (1983). A Boolean complete neural model of adaptive behavior. *Biological Cybernetics, 49*, 9–19.

Mitchell, T. M. (1982). Generalization as search. *Artificial Intelligence, 18*, 203–226.

Razran, G. (1965). Empirical codifications and specific theoretical implications of compound-stimulus conditioning: Perception. In W. F. Prokasy (Ed.), *Classical conditioning*. New York: Appleton-Century-Crofts.

Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology, 66*, 1–5.

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

Saavedra, M. A. (1975). Pavlovian compound conditioning in the rabbit. *Learning and Motivation, 6*, 314–326.

Schlimmer, J. C. (1986). *A comparison of feature salience measures* (Technical report #86-99). Irvine, California: The University of California, Department of Information and Computer Science.

Wasserman, E. A., Chatlosh, D. L., & Neunaber, D. J. (1983). Perception of causal relations in humans: factors affecting judgments of response-outcome contingencies under free-operant procedures. *Learning and Motivation, 14*, 406–432.

Whitlow, J. W., Jr., & Wagner, A. R. (1972). Negative patterning in classical conditioning: Summation of response tendencies to isolable and configural components. *Psychonomic Science, 27*, 299–301.

Woodbury, C. B. (1943). The learning of stimulus patterns by dogs. *J. comp. Psychol., 35*, 29–40.