STIMULUS CONTROL AND ASSOCIATIVE LEARNING

BEN A. WILLIAMS

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Interest in operant research on stimulus control has declined at the same time that much interest has burgeoned in nonoperant areas. Several examples of this shift toward traditional learning theory are considered, all of which have sponsored theoretical approaches that attempt to characterize the underlying associative units. These theoretical approaches are defended on the grounds that they have generated a deeper understanding of a variety of often puzzling phenomena. My projection is that future research will be determined even more strongly by theories about the structure of associations. Particular issues for which such discussion will have major impact include (1) whether conditional stimulus control is qualitatively different than simpler forms of stimulus control, (2) whether stimulus control is organized hierarchically, and (3) the origin of categories of stimulus equivalence.

Key words: stimulus control, associationism, conditional discrimination, expectancy, inhibition, concept learning, hierarchical organization

The prospects of operant research on stimulus control, if projected from trends reflected in the *Journal of the Experimental Analysis of Behavior* over the past 10 to 15 years, are diminishing rapidly. Perusal of research published in the 1960s reveals great interest in the topic, but this interest has decreased progressively to the point that few relevant papers are now published. Moreover, many of these are concerned with specialized issues such as the role of short-term memory or the relation between signal-detection theory and the matching law. Some part of this decline can be ascribed to several of the significant problems of the 1960s having been resolved (see reviews by Rilling, 1977, and Honig & Urcuioli, 1981). However, the burgeoning of research on related topics published in other journals (e.g., *Journal of Experimental Psychology: Animal Behavior Processes*) shows that this cannot be the entire explanation. Instead, what appears to have happened is that the interpretive framework provided by operant analyses in the 1960s

has been displaced by alternative theoretical approaches that have inspired investigations of stimulus control at different levels of analysis.

Contemporary research on stimulus control cannot be understood properly without some discussion of its two separate research traditions that remain distinct even today. The first is the "S-R" approach, originated by Thorndike but with most of its explanatory concepts drawn from Pavlov (1927), particularly as represented by the classic papers of Spence (1936, 1937). The cardinal assumption behind this tradition has been that the unit of learning is the association between stimulus and response. Consequently, stimulus control is not a topic in its own right, because stimulus control is nothing more than learning per se. In contrast is the position of Skinner (1938, 1953), which centers around the concept of the three-term contingency. Accordingly, the function of the stimulus is not to elicit a response directly, but rather to "set the occasion" for the response. Learning effects are differentiated into two components: that involving the response-reinforcer relation, and that involving how the stimulus comes to serve as the discriminative cue for control by the response-reinforcer relation.

This paper is dedicated to the memory of Don Hake. Preparation was supported by NIMH Research Grant R01 MH 35572-02. Requests for reprints should be addressed to the author, Department of Psychology, C-009, University of California, San Diego, La Jolla, California 92093.

Although these two approaches may seem to differ primarily in their specification of associative relations underlying stimulus control, that distinction is not the central focus of their critical differences. Instead, the difference is the status of "associationism" as a framework for understanding learning. Whereas the S-R tradition (and the "S-S" tradition) has sponsored theoretical speculations about the nature of the associative connections underlying observed behavioral changes, the Skinnerian tradition has remained antipathetic to any version of associationism (e.g., Skinner, 1977). Consequently, the concept of the three-term contingency has not been regarded as a unit of learning, but instead as a purely descriptive framework for describing the effects of stimulus-control procedures.

The thesis of the present paper is that progress toward understanding the processes of stimulus control has been impeded by the anti-theoretical posture of the Skinnerian tradition (e.g., Skinner, 1950), and that recent developments in the analysis of associative learning clearly demonstrate the importance of understanding it in terms of underlying associative units. Thus, future progress will go hand in hand with an adequate characterization of the different types of associations that may occur in a conditioning procedure, and of how they compete among themselves for acquisition and combine to generate behavior. To support this thesis, I will review particular examples of stimulus control that support the utility of associationistic analyses. I will then turn to several issues for future research, the resolution of which would constitute a major advance in our understanding of the fundamental processes underlying stimulus control.

THE CASE FOR ASSOCIATIONISM

Necessary Conditions for Stimulus Control

The one issue that dominated early operant research on stimulus control was the identification of discrimination contingencies necessary for a cue to be established as a discriminative stimulus. One position,

derived from S-R theory, was that any stimulus present at the time a response was reinforced would gain control. The second position, influenced strongly by the classic paper by Lashley and Wade (1946) criticizing S-R theory, was that the subject required discrimination training involving at least two values of the same stimulus dimension. In his influential review of research on this issue, Terrace (1966) argued that neither position was correct, because stimulus control did require some type of differential reinforcement with respect to the stimulus, but the differential reinforcement need not involve a discrimination between different points on the same dimension. Terrace's conclusion was of course heavily influenced by the important work of Jenkins and Harrison (1960), who demonstrated that presence/absence discrimination with respect to a 1000-Hz tone was sufficient to establish stimulus control by the tone-frequency dimension, but that simply reinforcing the bird's key pecking in the presence of the tone was not.

The problem with Terrace's conclusion is that it lacks generality in at least two respects. First, there are cases in which differential reinforcement of any sort appears unnecessary. For example, Hearst (1962) trained monkeys to press a bar in the presence of a houselight with constant illumination, and then tested for generalization to other light intensities. The result was strong control by the brightness dimension, despite the houselight not having been involved in any differential contingency of reinforcement. Of course, such an interpretation can be questioned either by appealing to implicit discrimination contingencies within the experiment (e.g., the monkey discriminates between the brightness within the chamber versus outside the chamber where the reinforcement contingency is inoperative), or by appealing to some type of experience prior to the experiment. Such appeals were in fact quite frequent during the 1960s, but we now know they are unnecessary.

A second limitation on the generality of Terrace's conclusion is that presence/absence

discrimination is sometimes not sufficient to produce stimulus control. This is illustrated in an experiment of Williams (1973), in which pigeons were trained to peck in the presence of one click rate, and not to peck in the absence of clicks. A second pair of pigeons were trained to peck during the same click rate, but not during a second click rate. The results, shown in Figure 1, were that control by the click-rate dimension occurred only for the subjects trained with the intradimensional discrimination.

Fig. 1. Generalization results taken from Williams (1973). Data in the upper panels are from subjects trained to discriminate a slow click rate as S+ from the absence of clicks as S-; data in the lower panels are from subjects trained to discriminate the same slow rate (#3) from a fast click rate (#8).

The reasons for the discrepancies between the various results are now understood. What was missing from Terrace's (1966) account was the notion of associative competition between stimuli. That is, whether a stimulus will gain control depends not only on the discrimination contingency correlated with that stimulus, but also on the other stimuli present that are potential competitors. One type of such competition is overshadowing that occurs when some other stimulus, equally correlated with the reinforcement contingency but with greater salience (as indexed by its faster rate of conditioning when both stimuli are presented

independently), is presented simultaneously with the target stimulus and competes with the target stimulus, preventing it from gaining control. An important example of this effect is provided by Rudolph and Van Houton (1977), who repeated the procedure of Jenkins and Harrison (1960) in which a tone was continuously presented while pigeons' key pecking was reinforced. Unlike Jenkins and Harrison, however, they trained some subjects to peck a darkened key. The result was that the subjects with the illuminated key failed to show control by the tone-frequency dimension, but those without key illumination demonstrated sharp control by the tone. Apparently, therefore, the reason for the initial failure of control was that the presence of the keylight overshadowed the tone for control.

Overshadowing presumably also occurred in the click-rate study shown in Figure 1. There, however, the source of overshadowing had to be an attribute of the click-rate stimulus itself. That is, the click-frequency dimension failed to gain control because some other attribute, presumably loudness, was more salient. The effect of the intradimensional discrimination training was then to eliminate control by this alternative dimension.

The mechanisms involved in overshadowing are still a subject of disagreement. One interpretation comes from the Rescorla-Wagner (1972) conditioning model that says, in essence, that an animal does not learn the relation between a new stimulus and reinforcement if the reinforcer is already predicted by other stimuli in the situation. In other words, the animals must be "surprised" by the reinforcer before any further learning involving the reinforcer will occur. A second interpretation comes from attention theory (e.g., Mackintosh, 1975) and says that a stimulus loses salience whenever it occurs conjointly with a second stimulus that better predicts the reinforcement. Still a third interpretation is that of Pearce and Hall (1980), who argue that a stimulus loses salience unless it is paired with events that are not already predicted. The empirical

issues separating these interpretations are complex and space does not permit their evaluation here. My purpose instead is to argue that none of them was likely to develop from operant research, because they are associative models that attempt to specify the structural features of stimulus relations that result from conditioning. In contrast, Skinner's concept of the three-term contingency was developed partly to avoid speculation about associations, and encouraged research to restrict itself to the level of determining correlations between discrimination-training conditions and stimulus-control outcomes.

Learning of "Expectancies"

A second example of the value of considering the nature of associative relationships comes from work usually done under the rubric of "expectancy" learning. Most of this research has used a conditional-discrimination task (sometimes referred to as symbolic matching to sample), which involves presenting a single sample cue (e.g., red or green) followed by a simultaneous choice between two other stimuli (e.g., a circle versus a line). The conditional-discrimination contingency is one whereby pecks on the circle are reinforced if red was presented as the sample and pecks on the line are reinforced if green was the sample. The finding of interest is that such learning is greatly facilitated if a different type of reward follows the two different correct responses. For example, the discrimination is learned more quickly if correct responses on the line are followed by food and correct responses on the circle are followed by water (Brodigan & Peterson, 1976).

Perhaps the most surprising finding is that such facilitation occurs when the different consequences following the two types of correct responses are different values of reward. For example, facilitation occurs if reward is presented immediately following one response but delayed for the second, compared to immediate reward for both responses (Carlson & Wielkiewicz, 1972). Most remarkable of all is the finding (Peterson, Wheeler, & Trapold, 1980) that facilitation

also occurs if one type of correct response is reinforced while the other is completely nonreinforced (except for advancing to the next trial). Such results violate our most basic notions about how reinforcement should work, because we usually assume that learning occurs more readily when behavior is reinforced than when it is not.

The interpretation given these findings is that animals learn to expect different trial outcomes as a function of the conditional cues that are presented, and that such expectancy learning occurs more quickly than learning which response goes with which sample stimulus (Peterson et al., 1980). Thus, the conditional cues become more distinctive because of their attached expectancies, and this increase in distinctiveness then facilitates the discrimination between the choice alternatives. Such an interpretation thus requires a distinction between different types of association, which seems unlikely to have been made if the analysis of stimulus control remained solely at the descriptive level of the stimulus "setting the occasion" for the response.

As a sidelight on this example, the use of distinctive types of reward seems to have considerable applied significance. For example, Litt and Schreibman (1981) trained autistic children to respond differentially to labels for different objects in their environment by picking up the object that had been named. Acquisition of the different labels was facilitated substantially when different reinforcers were contingent on correct behavior with respect to the different labels, rather than the same reinforcer being used for all correct labels.

Inhibition: An Internal State

A third example of the importance of understanding associative relationships concerns the concept of inhibition. As is well known, Skinner (1938) rejected the concept because he regarded it as indistinguishable from the notion that response strength is decreased when reinforcement is removed. As is also well known, history has not dealt well with that aspect of Skinner's position;

there is now convincing evidence for the reality of inhibition as a behavioral process (see Hearst, Besley, & Farthing, 1970). Perhaps the most convincing example is the "superconditioning" effect in respondent conditioning that is obtained by first establishing some stimulus (X) as a cue for nonreinforcement (e.g., $A = S+$, $A + X = S-$), and then presenting that cue with a new stimulus (B), with the compound ($B + X$) then paired with the reinforcer (see Rescorla, 1971). The critical finding is that conditioning to the new stimulus occurred more readily when it was conjoined with the conditioned inhibitor than when it was paired with the reinforcer alone (the opposite of overshadowing). Such a finding implies that the stimulus trained as a conditioned inhibitor is not neutral, nor simply associated with the absence of reinforcement, but instead has active properties resulting from its status as a cue for nonreinforcement.

Despite the clear demonstration of conditioned inhibition, our understanding of the associative relations involved is far from complete. Several different characterizations of what is learned in a conditioned inhibition procedure have been offered (see Dickinson, 1980, pp. 95-101). In order of decreasing specificity: (1) the stimulus is associated with the absence of the particular reinforcer used in training, and so will inhibit any behavior maintained by that reinforcer; (2) the stimulus acquires inhibitory properties with respect to the particular motivational system activated at the time of acquisition, and thus will affect all behavior pertaining to that motivational system but not behavior pertaining to other motivational systems; (3) the stimulus acquires general inhibitory properties, and thus will reduce any operant behavior, however that behavior is maintained.

An example of research relevant to these predictions is provided by Cotton, Goodall, and Mackintosh (1982), who trained rats on a conditioned-suppression procedure in which conditioned inhibition was established by the stimulus being correlated with a reduction in the level of shock intensity.

Two different stimuli (A and B) initially were paired separately with a 1.0-mA shock; then conditioned inhibitory training occurred in which one of these stimuli was presented in compound with the stimulus (X) targeted as a conditioned inhibitor, with this compound ($A + X$) paired with a 0.4-mA shock. The inhibitory property of the target stimulus was then demonstrated by pairing it with the second stimulus used in the initial training ($B + X$), with the result that the level of suppression to that stimulus was reduced. The generality of the inhibitory effect was investigated in a second experiment involving the same acquisition procedure (A paired with a 1.0-mA shock, $A + X$ paired with a 0.4-mA shock), but with the test for conditioned inhibition being the effects of stimulus X when compounded with a second stimulus that initially was paired with a 0.4-mA shock itself. Suppression was not reduced, implying that the effects of the conditioned inhibitor were specific to the particular level of shock intensity used during the establishment of the conditioned inhibitor.

Interpretation of the specificity of conditioned inhibition is complicated by still other data showing that a stimulus trained as a conditioned inhibitor with respect to shock may retard subsequent conditioning involving food as the reinforcer. Fowler, Goodman, and Zanich (1977) established white noise as a conditioned inhibitor with respect to shock in a Pavlovian discrimination procedure, and then presented the noise contingent on the correct response in a T maze. Correct choices were also followed by food, with the noise intervening between the choice point and entry into the goal box. The result was that T-maze acquisition was retarded, relative to a second group that received the noise randomly related to shock during the initial training. Thus, a conditioned inhibitor with respect to shock "blocked" acquisition of the food-reinforced response. In contrast, prior pairing of the noise with shock, to establish it as a conditioned excitor with respect to shock, was shown to facilitate acquisition of the T-maze discrimination, an

effect similar to the “superconditioning” effect noted earlier. Such results suggest that conditioned inhibition with respect to aversive stimuli is functionally similar to conditioned excitation with respect to appetitive reinforcers, inasmuch as conditioned excitors with respect to shock act as conditioned inhibitors with respect to food, and vice-versa. Such cross-motivational interactions suggest that conditioned-inhibition effects may be quite general in their ramifications. Whatever the eventual outcome of this issue, its lesson for the present discussion is that the nature of the preexisting associations between the elements of the conditioning situation must be understood before the effects of particular conditioning operations can be predicted.

FUTURE DEVELOPMENTS

The Interpretive Status of Conditional Relations

The preceding discussion has been critical of the “operant” treatment of stimulus control because it has eschewed consideration of the units of learning in terms of “what is learned.” I now will argue that the most fundamental operant concept, the three-term contingency, is central to future progress in the analysis of stimulus control. Whereas traditional associatioinistic accounts (e.g., S-R and S-S theories) have attempted to reduce conditioning effects to some set of binary associations (see Hearst, 1975, for a modern example), the notion that a stimulus “sets the occasion” for the response implies that such binary associations provide an inadequate account because stimulus control is

inherently conditional in nature (see Jenkins, 1977, for a related discussion). The issue becomes, therefore, whether conditional relations are of a fundamentally different sort than simple associations, and if so, how such different relations coexist, and interact, in the organization of behavior.

To see the importance of distinguishing among different types of stimulus control, consider the issue of conditioned reinforcement. As is well known, a long-standing claim has been that the necessary and sufficient condition for a stimulus to become a conditioned reinforcer is that it serve as a discriminative stimulus. But perhaps the conditioned reinforcement properties may depend on exactly what type of discriminative function the stimulus serves. Shown in Table 1 are three conditions, in all of which the stimulus is in the same temporal relation to reinforcement. In Condition 1, a keylight is presented in a discrete-trial procedure, with trials terminating in a 50% food-delivery schedule independent of responding (i.e., an autoshaping procedure). In Condition 2, the same keylight is presented in conjunction with a second keylight in a discrete-trial procedure involving two types of trials, those in which the white keylight is presented alone and no reinforcement is available, and those in which the red and white are presented together with pecks on white reinforced on 50% of the trials. For Condition 3, the red keylight is presented in conjunction with both white and blue keylights on the second key. In addition, the first key also has presentations of a green keylight, which also is presented in conjunction with both the

Table 1
An outline of training conditions for conditioned reinforcement. Each row represents a different type of trial, including the stimuli and reinforcement contingencies. The same procedures constitute pretraining for the training procedures shown in Table 2.

Condition 1	RED:	50% food
Condition 2	RED + WHITE: WHITE alone:	50% food for pecks to white No food
Condition 3	RED + WHITE: RED + BLUE: GREEN + BLUE: GREEN + WHITE:	100% food for pecks to white No food 100% food for pecks to blue No food

white and blue lights on the second key. The critical feature is the presence of a conditional discrimination contingency such that pecks on white are reinforced only in conjunction with red, and pecks on blue are reinforced only in conjunction with green. Thus, combinations of red + blue and green + white are not reinforced.

The issue is the value of the red keylight in all three conditions, as assessed by a conventional conditioned reinforcement assay. The three are similar in that presentations of red are correlated with a 50% reinforcement schedule. Despite this similarity, intuition suggests a considerable difference in the information conveyed in the three conditions. To the extent that such differences correspond to different conditioned reinforcement values, previous analyses of stimulus control have no easy way of accounting for them.

A related example involves the phenomenon of blocking (see Kamin, 1969). Recall that a blocking procedure involves pretraining in which one stimulus element is paired with the reinforcer and then presented in compound with a new element, with this compound paired with the same reinforcer. The result is less conditioning to the new element than would have occurred if the pretraining had not occurred. The issue for the present discussion is what types of blocking

effect would be produced by the red keylight used in the procedures shown in Table 1. Shown in Table 2 is a set of procedures that provide paradigms for blocking in which the red keylight, with versus without the pretraining shown in Table 1, is paired with some new stimulus. It should be evident that the three procedures for assessing blocking effects (Table 2) are analogous to the three pretraining procedures (Table 1), as one involves autoshaping, the second a discriminated operant procedure, and the third a true conditional discrimination. At issue is whether the specific discriminative function established during pretraining generates blocking effects only when the same specific discriminative functions are involved during the second phase of training. If such specificity does occur, a strong case is made for distinguishing among different kinds of associative relations, inasmuch as competition for stimulus control would occur only when the alternative stimuli compete for the same discriminative function.

A second approach to differentiating various types of associative relations involves "learning to learn." The issue is whether the improvement in learning proficiency seen in such procedures (learning set, serial reversal learning) is specific to particular types of associative relations. For example, can con-

Table 2

Training procedures for different blocking experiments. For each experiment four groups of subjects would be compared: three with the pretraining conditions shown in Table 1, a fourth with no prior training with the red stimulus.

	<i>Training</i>	<i>Test</i>
EXPERIMENT 1:	RED + CIRCLE: 50% food	CIRCLE-alone
EXPERIMENT 2:	(RED + TONE) + CIRCLE: 50% food for CIRCLE pecks CIRCLE-alone: No food	TONE + CIRCLE versus CIRCLE-alone
EXPERIMENT 3:	(RED + TONE) + CIRCLE: 100% food for CIRCLE pecks (RED + TONE) + LINE: No food (GREEN + NOISE) + LINE: 100% food for LINE pecks (GREEN + NOISE) + CIRCLE: No food	TONE + CIRCLE versus TONE + LINE versus NOISE + LINE versus NOISE + CIRCLE

ditional discriminations be facilitated independently of simple discrimination learning? If so, it would suggest that an animal may be sensitized to the conditional relation *per se*, in a manner analogous to training the animal to attend to simple stimulus attributes such as hue or brightness.

An example of how such sensitization can be demonstrated is provided by a previously unpublished study from my laboratory using a serial-reversal procedure. The design of the study is shown in Table 3. Four separate groups of four subjects were given two phases of training, the second of which was common to all groups (as shown on the right side of the table). The procedure for the second phase was a conditional discrimination in which the relation between the conditional cues and the choice cues was reversed after each separate problem was learned to criterion. On any particular trial, red or blue appeared as the conditional cue; the choice cues were simultaneous presentations (on separate keys) of a blue 150° line and a yellow 120° line. Thus, on the first problem, the 150° line was correct if red had served as the conditional cue and choice of the 120° line was correct if blue had been the conditional cue. Then on the second problem, the 120° line was correct with red as the conditional cue, and choice of the 150° line was correct with blue as the conditional cue. Each problem was terminated after 10 consecutively correct trials.

At issue was whether learning the reversals of the conditional discrimination in the second phase of training would be facilitated selectively by different types of training in Phase 1. Groups CD-W/G and CD-R/B both received reversal training with a different conditional discrimination in which the choice cues were a 0° green line and a 45° red line. The two groups differed with respect to the conditional cues, as the latter received the same conditional cues as were used in Phase 2 (red and blue), whereas the former received different conditional cues (white and green). The two remaining groups were trained on a serial reversal problem of a simple discrimination in which particular choice cues (the green 0° line and the red 45° line) were correct or incorrect throughout a session. Reversals of the reward values of the choice cues were then instituted when each successive problem had been learned to a criterion of 10 consecutively correct trials. In addition, all subjects in the serial reversal groups could produce illumination of the choice cues only by pecking at the colors used as conditional cues. Unlike the conditional-discrimination groups, however, there was no consistent relation between the color that appeared at the start of the trial and the reward value of the choice cues on that particular trial. The serial-reversal groups were included in order to control for any generalized effects of extended discrimination training, and in order to equate

Table 3
Design for "Learning to Learn" Experiment Involving a Conditional Relation

Group	Phase 1		Phase 2	
	Conditional/Observing Cues	Choice Cues	Conditional Cues	Choice
CD-W/G	white/green	green 0° red 45° lines	red/blue	blue 150° yellow 120° lines
CD-R/B	red/blue	green 0° red 45° lines	red/blue	blue 150° yellow 120° lines
SR-W/G	white/green	green 0° red 45° lines	red/blue	blue 150° yellow 120° lines
SR-R/B	red/blue	green 0° red 45° lines	red/blue	blue 150° yellow 120° lines

exposure to the various discriminative stimuli.

The results for the various groups are shown in Figure 2, in terms of the average number of errors required to learn each new reversal. The primary results of interest are those of Phase 2, during which all groups received training on serial reversals of the same conditional discrimination problem. In general, all groups exhibited the same pattern of performance—a large increase in the number of errors required to learn the first few reversals and then a substantial improvement. Most importantly, three of the four

groups were also similar in the absolute number of errors, the one exception being Group CD-R/G, which was trained with the same conditional cues used during Phase 1.

Two aspects of the results shown in Figure 2 are relevant to the issue of whether subjects can be sensitized to learning conditional relationships. Considering first the results for Group CD-W/G: Its failure to perform better than the two serial-reversal groups shows that conditional discrimination training on one set of stimuli did not produce a generalized improvement in conditional discrimination learning involving a completely different set of stimuli. Such a finding thus argues against “abstraction” of the conditional relation per se. However, the facilitation that did occur for Group CD-R/B suggests that sensitization to the conditional relation may occur, but only with respect to the particular stimuli involved in prior training. Such a restriction is perhaps not surprising, given that different types of relations may be relevant depending upon which stimuli are present. Thus, the particular stimuli serve to cue the particular relation that is relevant, in a manner analogous to their cueing the correct choice stimulus. The results shown in Figure 2 are not definitive evidence in favor of this interpretation, but they do provide encouragement for attempts to isolate the specific associative relationships underlying the learning-to-learn effect.

The notion that a stimulus instructs the subject that a conditional relation is in effect has been anticipated by previous investigators. For example, in their review of matching and oddity experiments, Carter and Werner (1978) discuss such learning in terms of a set of “if . . . then . . .” rules. Similarly, Cumming and Berryman (1965) interpreted such learning in terms of the sample stimulus serving an “instructional” function. In neither case, however, was the status of the conditional relation discussed directly. The purpose of the present discussion is to explicitly encourage such discussion, with the eventual aim of classifying the different types of stimulus-control relations that may be permitted by any conditioning

Fig. 2. Discrimination performance for the two phases of study of serial-reversal training of a conditional discrimination. The first segment shows the results of Phase 1, in which the different groups received different discriminations. Note that the groups designated by circles were trained on a conditional discrimination; the groups designated by triangles were trained on serial reversals of a simple simultaneous discrimination (see Table 3). The second segment shows the results of Phase 2, in which all groups received the same conditional discrimination. Twelve reversals were presented in Phase 1, eight reversals in Phase 2. The ordinate refers to the geometric mean of each group.

situation, and of characterizing how the different relationships interact, and possibly compete, for control of behavior.

Hierarchical Levels of Stimulus Control

As noted in the discussion of the conditions necessary to establish stimulus control, the concept of associative competition has provided the key ingredient to most theoretical analyses of the past 10 to 15 years. The most widely accepted account of such competition effects is the Rescorla-Wagner (1972) model of classical conditioning, which postulates that the fundamental determinant of conditioning is the discrepancy between the existing associative strength and the maximum possible given the reinforcer. As is well known, their model accounts for a wide variety of conditioning effects, including some that are otherwise paradoxical (e.g., see Kamin & Gaioni, 1974).

The usual interpretation of associative competition effects has been challenged by recent studies in which "reminder" treatments reinstated responding to the stimulus that supposedly was not associated with the reinforcer (Balaz, Gutsin, Cacheiro, & Miller, 1982; Schachtman, Gee, Kaspro, & Miller, 1983). After the usual blocking procedure in which the blocked stimulus was shown not to evoke responding, additional, unreinforced exposure to the unconditioned stimulus, the blocked stimulus, or the conditioning context resulted in the emergence of the conditioned response to the blocked stimulus, despite its not having been paired with the reinforcer outside the blocking procedure. Such results have been interpreted as showing that the deficits produced by blocking procedures are due not to the failure to establish associative connections, but rather to prevention of their subsequent "retrieval."

But the retrieval interpretation is itself challenged by other results. For example, in a modification of the usual blocking procedure, Kremer (1978) pretrained his subjects with two stimuli, A and B, each separately paired with the reinforcer until an asymptotic level of conditioning had occurred. The

A and B elements were then compounded, a third stimulus (C) was added, and the pairing of the ABC compound with the reinforcer was continued. Later the C element was tested in isolation and found to be a conditioned inhibitor. The apparent reason is that separate pretraining with the A and B elements caused the existing associative strength of the ABC compound to be greater than the asymptotic level possible, and this negative discrepancy was associated with the addition of a new conditioning element (C), thus causing that new element to become inhibitory. The important point for the present discussion is that such inhibitory learning should not have occurred if the effect of a blocking procedure is only to prevent later retrieval of conditioning to the new element. The issue, therefore, is what variables determine when retrieval effects will play a role, as opposed to blocking being mediated by competition between the associative relationships themselves.

A related puzzle is how the blocking phenomenon is to be reconciled with the recent conditioning model of Gibbon and his collaborators—the scalar expectancy theory (Gibbon, 1981; Gibbon & Balsam, 1981). Their model provides an accurate account of a variety of conditioning effects in classical conditioning, including the effects of inter-trial interval, trial duration, and percentage reinforcement. Its fundamental tenet is that conditioning is to be understood in terms of the ratio of the rates of reinforcement in the presence of the CS to that in the situation as a whole. The subject is assumed to form independent estimates of the reinforcement rates during the CS and during the situation, and when that ratio exceeds some critical value, the conditioned response emerges. The problem is that the assumption of independent learning about the two kinds of reinforcement rates conflicts directly with the concept of blocking, which posits associative competition between the two types of learning. As noted earlier, there is strong evidence supporting the blocking concept, so the issue again is how the conflict is to be resolved.

An outline for the possible resolution of these issues comes from the recent suggestion of Balsam (in press) that stimulus control may be organized hierarchically. Thus, within each level of the hierarchy there may be associative competition between the various stimuli that predict reinforcement, whereas no such competition occurs between levels. Instead, the stimuli at the higher levels "set the occasion" for the reinforcement contingencies that lie below them. In addition, a given stimulus might serve both discriminative functions simultaneously, both being itself in direct association with reinforcement and setting the occasion for stimulus control by other cues in the situation.

As an example of the application of the notion of hierarchical organization, consider the provocative findings of Lindblom & Jenkins (1981). They first established key pecking with the usual autoshaping procedure and then investigated the effects of adding extra reinforcers during the intertrial interval (ITI)—a procedure similar to the random control procedure introduced by Rescorla (1967). Like previous findings, the effect of their added reinforcers was to abolish the conditioned response, presumably because the keylight no longer was predictive of the occurrence of reinforcement. However, Lindblom and Jenkins then removed all reinforcement from the situation, both following the keylight and during the ITI, and found that the conditioned response emerged in a strength comparable to that before the extra reinforcers were added. Thus, the change in contingency did not abolish the association between the keylight and food, but instead only suspended the ability of the keylight to generate responding. In a second experiment, however, they presented the random-control procedure from the start of training, without first establishing the keypeck response, and found that their extinction procedure had no effect (i.e., no responding occurred after extinction). In the latter case, therefore, the effect of the extra reinforcers presumably was to prevent the association between key peck-

ing and food from being established.

Given the different effects of the random-control procedure as a function of whether key pecking was already established, a possible interpretation is that the extra reinforcers presented in the ITI served a different discriminative function in the two cases. When presented from the start of training, such reinforcers presumably were associated with the background cues, thus causing the background cues to block conditioning to the keylight (Rescorla & Wagner, 1972). But when presented later in training, the extra reinforcers assumed the status of discriminative stimuli, thus setting the occasion for the keylight-food contingency no longer being in effect. Moreover, the status of the extra reinforcers as discriminative stimuli then prevented their association with the background cues, thus preventing the blocking effects that usually would be expected (see Tomie, 1981).

The foregoing explanation is of course highly speculative. The important point is that the notion of hierarchical organization allows stimuli to participate in several different types of associative relationships, so that associative competition may or may not occur, depending upon the nature of those relationships. The fact that associative competition is a critical feature of conditioning necessitates the consideration of such issues; it implies that any analysis of stimulus control must go beyond the reinforcement contingencies correlated with stimuli in isolation. Such analyses must include, in addition, the nature of other potential stimulus-reinforcer relations in the same situation, and, on the basis of the preceding discussion, the nature of possible higher-order relationships as well.

Origins of Stimulus Categories

A final issue to be considered is control by stimulus categories, such that discriminative training on only a few exemplars of such categories produces generalization to other, novel, instances. Such generalization has been demonstrated with a variety of stimulus classes, including people, trees, fish, bodies

of water, pigeons, pictures of a particular person, leaves, and man-made objects (Herrnstein, 1979). Despite such frequent demonstrations, however, relatively little analysis has been directed toward the origin of such categories or to how, if at all, it is related to stimulus control of a simpler variety.

At least two separate theories of such learning have been offered, both of which come from experiments with human subjects. One, known as prototype theory (Posner & Keele, 1968), assumes that experience with different exemplars produces an "abstraction" of the central tendency of the category, so that future instances are simply compared with this ideal type. A second view (known as feature theory), much closer to the concept of stimulus generalization, is that particular exemplars each possess several different features, each of which acquires response strength in its own right (Hayes-Roth & Hayes-Roth, 1977). Thus, new exemplars are classified as instances of the concept if the sums of their features are sufficiently close to the features shared by the training exemplars. Evidence separating these accounts has been ambiguous (e.g., Elio & Anderson, 1981).

A precursor to the analysis of category learning is the specification of the nature of the features that comprise a conceptual instance, and the nature of the combination rules by which different features exert their effects. One possible rule is simple additive summation, which implies that each feature contributes to the probability of responding according to its salience and previous history of reinforcement. Thus, when the sum of these probabilities exceeds some threshold, the classification response occurs. Such a view implies that category learning is no different than simple stimulus control, except for the nature and number of features that are involved. An alternative view is that the different features interact, such that various conjunctions and disjunctions control response probabilities different from the sum of response strengths controlled by their constituent elements. Such combinations might then be regarded as "features" in their own

right. It should be apparent that such interactive combination rules are simply a variety of conditional stimulus control, so that the likelihood that such rules would function should vary with the salience of the conditional relation as a discriminative event. The interactive combination rule also seems to imply, however, that stimulus features that combine by a conditional rule may assume some type of configural integrity. Understanding the emergence of such unified stimulus compounds may thus constitute a critical part of understanding the nature of concept learning.

Previous work on the acquisition of natural concepts in animals in fact reveals very little about whether such concepts represent distinct conceptual classes. Correct classification of novel instances says nothing about the underlying basis, as the category may be mediated by a single concept, or by several functionally independent concepts, each of which is reducible to some combination of perceptual features. Herrnstein, Loveland, and Cable (1976) and Herrnstein (1979) have argued against a single prototype as underlying the various natural categories and have argued instead for Wittgenstein's notion of "family resemblances," whereby a conceptual class can be regarded as an ill-defined region in some multidimensional similarity space. Such a conceptualization seems similar to the feature theory noted earlier, but Herrnstein et al. provide no account of how the dimensions of the similarity space are generated.

A fundamental and neglected issue regarding the acquisition of natural concepts is whether stimulus categories can be "induced." That is, do categories depend upon innately given dimensions of similarity, so that the process of acquisition is simply learning to attend to the relevant features (acknowledging that the features themselves may be quite abstract), or is it possible instead to cause two independent stimuli to be regarded as instances of the same class despite their sharing no physical similarity?

One method by which such induction might occur is suggested by recent studies of

classical conditioning showing that features of a stimulus become associated simply by virtue of their simultaneous pairing. For example, a recent controversy in taste-aversion learning has been over the question of why overshadowing does not occur when the stimulus compound consists of taste + odor. That is, pairing a taste + odor compound with poison does not reduce the amount of conditioning to the odor alone, but instead potentiates it, compared to that obtained when the odor alone is paired with the poison (Rusiniak, Hawkins, Garcia, & Brett, 1979). The apparent reason for this effect is that simultaneous pairing of the taste and odor causes the two to become associated, so that whichever associations occur to one will be acquired by the other (Durlach & Rescorla, 1980). Such an effect is of course similar to "sensory preconditioning," but unlike the investigation of that earlier concept, recent evidence has shown such effects to be a major determinant of the outcomes of many different conditioning procedures (Rescorla, 1981).

To see how mediating associations might play a role in concept acquisition, consider a hypothetical case in which each exemplar of the concept has only two different features, and no two exemplars share the same two features. Thus, the features comprising three different instances might be designated AB, BC, and CD. The issue is how the first (AB) and third (CD) come to be regarded as instances of the same class despite having no features in common. In fact, categorization follows directly from the notion of mediating associations, because features B and D are both associated with C, and thus should become associated with each other. Such mediation would quickly generate the type of similarity space seemingly implied by the notion of "family resemblances." As yet, however, the potential role of such mediating associations has received essentially no analysis.

It is possible that the abstraction of stimulus categories depends upon behavioral processes not reducible to families of associations. Recent investigations of stimulus equivalence by Sidman and his collaborators

(Sidman, Rauzin, Lazar, Cunningham, Tailby, & Carrigan, 1982; Sidman & Tailby, 1982) support this possibility, as they demonstrated that equivalence between arbitrarily related stimuli could be developed with human subjects but not with other primates. Their various subjects were trained on a "symbolic" matching-to-sample procedure in which stimuli in Class A were first used as the sample stimuli for one set of choice stimuli (Set B), and then again as the sample stimuli for a second set of choice stimuli (Set C). They then determined whether the equivalence established between A and B and between A and C would generate equivalence between B and C. Such equivalence did occur for young children, but not with either monkeys or baboons. Sidman et al. argued against the presence of verbal labels as the cause of the species differences, but what did produce the differences remains uncertain.

The research by Sidman and his collaborators is virtually unique in its investigation of how arbitrary conceptual categories may be created. Given that the origin of stimulus equivalence is perhaps the most venerable issue in the study of cognition, the applicability of behavior analyses to that issue, with both human and nonhuman subjects, will have major implications for the future of research on complex stimulus control.

CONCLUSIONS

The preceding three examples are only a small portion of the avenues that future research on stimulus control will take. I have discussed a few specific examples extensively, rather than survey a wide panorama, partly because they now appear to be the most important issues for the future and partly because some level of detail is necessary to exemplify the kinds of analyses that appear to be required. The examples chosen also share a common thesis: Research on stimulus control has reached the point at which consideration of the "units of learning" has become a critical issue. Such consideration represents a notable departure from the

entirely descriptive analyses that have dominated past research, and will be regarded by many as anathema because of its excursion into the murky waters of hypothetico-deductive theorizing. The present discussion has attempted to defuse such objections by demonstrating the utility of such constructs as "associative competition," "mediated associations," "response-outcome expectancies," and the like. Whatever the particular merits of such concepts, recent research has shown clearly that some consideration is required of the "state" of the organism at the moment that learning experiences occur, because very different outcomes may result from identical procedures depending upon the subject's previous history. Speculations about the internal processes mediating such history effects have led to important recent advances and should be a major component of future accounts of stimulus control.

REFERENCES

- Balaz, M. A., Gutsin, R., Cacheiro, H., & Miller, R. R. (1982). Blocking as a retrieval failure: Reactivation of associations to a blocked stimulus. *Quarterly Journal of Experimental Psychology*, **34B**, 99-113.
- Balsam, R. (in press). Bringing the background to the foreground: The role of contextual cues in autoshaping. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Vol. 4. Discrimination processes*. Cambridge, MA: Ballinger.
- Brodigan, D. L., & Peterson, G. B. (1976). Two-choice conditional discrimination performance of pigeons as a function of reward expectancy, prechoice delay, and domesticity. *Animal Learning & Behavior*, **4**, 121-124.
- Carlson, J. G., & Wielkiewicz, R. M. (1972). Delay of reinforcement in instrumental discrimination learning of rats. *Journal of Comparative and Physiological Psychology*, **81**, 365-370.
- Carter, D. E., & Werner, T. J. (1978). Complex learning and information processing by pigeons: A critical analysis. *Journal of the Experimental Analysis of Behavior*, **29**, 565-601.
- Cotton, M. M., Goodall, C., & Mackintosh, N. J. (1982). Inhibitory conditioning resulting from a reduction in the magnitude of reinforcement. *Quarterly Journal of Experimental Psychology*, **34B**, 163-180.
- Cumming, W. W., & Berryman, R. (1965). The complex discriminated operant: Studies of matching-to-sample and related problems. In D. I. Mostofsky (Ed.), *Stimulus generalization* (pp. 284-330). Stanford, CA: Stanford University Press.
- Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.
- Durlach, P. J., & Rescorla, R. A. (1980). Potentiation rather than overshadowing in flavor-aversion learning: An analysis in terms of within-compound associations. *Journal of Experimental Psychology: Animal Behavior Processes*, **6**, 175-187.
- Elio, R., & Anderson, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory*, **7**, 397-417.
- Fowler, H., Goodman, J. H., & Zanich, M. L. (1977). Pavlovian aversive to instrumental appetitive transfer: Evidence for across-reinforcement blocking effects. *Animal Learning & Behavior*, **5**, 129-134.
- Gibbon, J. (1981). The contingency problem in autoshaping. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 285-308). New York: Academic Press.
- Gibbon, J., & Balsam, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (219-253). New York: Academic Press.
- Hayes-Roth, B., & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning & Verbal Behavior*, **16**, 321-338.
- Hearst, E. (1962). Concurrent generalization gradients for food-controlled and shock-controlled behavior. *Journal of the Experimental Analysis of Behavior*, **5**, 19-31.
- Hearst, E. (1975). The classical-instrumental distinction: Reflexes, voluntary behavior, and categories of associative learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 2. Conditioning and behavior theory* (pp. 181-223). Hillsdale, NJ: Erlbaum.
- Hearst, E., Besley, S., & Farthing, W. (1970). Inhibition and the stimulus control of operant behavior. *Journal of the Experimental Analysis of Behavior*, **14**, 373-409.
- Herrnstein, R. J. (1979). Acquisition, generalization, and discrimination reversal of a natural concept. *Journal of Experimental Psychology: Animal Behavior Processes*, **5**, 116-129.
- Herrnstein, R. J., Loveland, D. H., & Cable, C. (1976). Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **2**, 285-302.
- Honig, W. K., & Urcuioli, P. J. (1981). The legacy of Guttman and Kalish (1956): 25 years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior*, **36**, 405-445.
- Jenkins, H. M. (1977). Sensitivity of different response systems to stimulus-reinforcer and response-reinforcer relations. In H. Davis & H. M. B. Hurwitz (Eds.), *Operant-Pavlovian interactions* (pp. 47-66). Hillsdale, NJ: Erlbaum.
- Jenkins, H. M., & Harrison, R. H. (1960). Effect of discrimination training on auditory generalization. *Journal of Experimental Psychology*, **59**, 246-253.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York: Appleton-Century-Crofts.

- Kamin, L. J., & Gaioni, S. J. (1974). Compound conditioned emotional response conditioning with differentially salient elements in rats. *Journal of Comparative and Physiological Psychology*, *87*, 591-597.
- Kremer, E. F. (1978). The Rescorla-Wagner model: Losses in associative strength in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 22-36.
- Lashley, K. S., & Wade, M. (1946). The Pavlovian theory of generalization. *Psychological Review*, *53*, 72-87.
- Lindblom, L. L., & Jenkins, H. M. (1981). Responses eliminated by noncontingent or negatively contingent reinforcement recover in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 175-190.
- Litt, M. D., & Schriebman, L. (1981). Stimulus-specific reinforcement in the acquisition of receptive labels by autistic children. *Analysis and Intervention in Developmental Disabilities*, *1*, 171-186.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276-298.
- Pavlov, I. P. (1927). *Conditioned reflexes* (G. V. Anrep, Trans. & Ed.). London: 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.
- Peterson, G. B., Wheeler, R. L., & Trapold, M. A. (1980). Enhancement of pigeons' conditional discrimination performance by expectancies of reinforcement and nonreinforcement. *Animal Learning & Behavior*, *8*, 22-30.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353-363.
- Rescorla, R. A. (1967). Pavlovian conditioning and its proper control procedures. *Psychological Review*, *74*, 71-80.
- Rescorla, R. A. (1971). Variation in the effectiveness of reinforcement and nonreinforcement following prior inhibitory conditioning. *Learning and Motivation*, *2*, 113-123.
- Rescorla, R. A. (1981). Simultaneous associations. In P. Harzem & M. D. Zeiler (Eds.), *Advances in analysis of behaviour: Vol. 2. Predictability, correlation, and contiguity* (pp. 47-80). Chichester, England: Wiley.
- 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 research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- Rilling, M. (1977). Stimulus control and inhibitory processes. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 432-480). Englewood Cliffs, NJ: Prentice-Hall.
- Rudolph, R. L., & Van Houten, R. (1977). Auditory stimulus control in pigeons: Jenkins and Harrison (1960) revisited. *Journal of the Experimental Analysis of Behavior*, *27*, 327-330.
- Rusiniak, K. W., Hawkins, W. G., Garcia, J., & Brett, L. P. (1979). Flavor-illness aversions: Potentiation of odor by taste in rats. *Behavioral and Neural Biology*, *25*, 1-17.
- Schachtman, T. R., Gee, J. L., Kasprow, W. J., & Miller, R. R. (1983). Reminder-induced recovery from blocking as a function of the number of compound trials. *Learning and Motivation*, *14*, 154-164.
- Sidman, M., & Tailby, W. (1982). Conditional discrimination vs. matching to sample: An expansion of the testing paradigm. *Journal of the Experimental Analysis of Behavior*, *37*, 5-22.
- Sidman, M., Rauzin, R., Lazar, R., Cunningham, S., Tailby, W., & Carrigan, P. (1982). A search for symmetry in the conditional discriminations of rhesus monkeys, baboons, and children. *Journal of the Experimental Analysis of Behavior*, *37*, 23-44.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review*, *57*, 193-216.
- Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. (1977). Why I am not a cognitive psychologist. *Behaviorism*, *5*(2), 1-10.
- Spence, K. W. (1936). The nature of discrimination learning in animals. *Psychological Review*, *43*, 427-449.
- Spence, K. W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review*, *44*, 430-444.
- Terrace, H. S. (1966). Stimulus control. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application* (pp. 271-344). New York: Appleton-Century-Crofts.
- Tomie, A. (1981). Effects of unpredictable food on the subsequent acquisition of autoshaping: Analysis of the context-blocking hypothesis. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 181-215). New York: Academic Press.
- Williams, B. A. (1973). The failure of stimulus control after presence-absence discrimination of click-rate. *Journal of the Experimental Analysis of Behavior*, *20*, 23-27.