JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1981, 36, 405-445 NUMBER 3 (NOVEMBER)

THE LEGACY OF GUTTMAN AND KALISH (1956): 25 YEARS OF RESEARCH ON STIMULUS GENERALIZATION

WERNER K. HONIG AND PETER J. URCUIOLI

DALHOUSIE UNIVERSITY

This paper is a selective review of the methods, problems, and findings in the area of operant stimulus generalization over the 25 years since the publication of the original paper by Guttman and Kalish (1956) on discriminability and spectral generalization in the pigeon. The paper falls into five main sections, which encompass the main themes and problems stemming from the Guttman and Kalish work and its immediate successors. The first section addresses the relationship between stimulus generalization and stimulus control, as well as the variety of testing procedures and dependent variables used to measure generalization. The next section reviews the limited literature on the effects of early rearing on the generalization gradient. The relationship between discriminability among test stimuli and the slope of the spectral gradient is discussed in the third section, with emphasis upon recent reassessments of the pigeon's hue discriminability function. The fourth section reviews the topic of inhibitory stimulus control, one which developed with the discovery of the peak shift following intradimensional discrimination training. Problems of definition and measurement are discussed in conjunction with the gradient forms used to index inhibitory control. The last section is devoted to attentional effects and the two principal theories postulated to account for them. A survey of different attentional paradigms is provided and the possible role of constant irrelevant stimuli as a source of control is examined. A brief conclusion summarizes the contribution of the generalization technique toward an understanding of the nature and acquisition of stimulus control.

Key words: Stimulus generalization, generalization gradients, early rearing, discriminability, inhibitory stimulus control, attention, pigeons

---

TABLE OF CONTENTS

INTRODUCTION

In 1956, Norman Guttman and Harry Kalish published an article in the *Journal of Experimental Psychology* reporting a study of stimulus generalization in pigeons. The design of their experiment was remarkably simple. First, they trained separate groups of pigeons to peck a key lighted by one of four spectral stimuli. A steady rate of responding was established by intermittently reinforcing pecking on a variable interval (VI) schedule. Guttman and Kalish then tested for stimulus generalization by repeatedly presenting 11 different spectral values in random order during extinction. The test values were generally spaced 10 nm apart, covered a range of 120 nm, and included the original training value. When the total number of test responses was plotted as a function of wavelength, orderly decremental gradients of generalization were revealed around each training value. The average gradients for the four groups are shown in Figure 1.

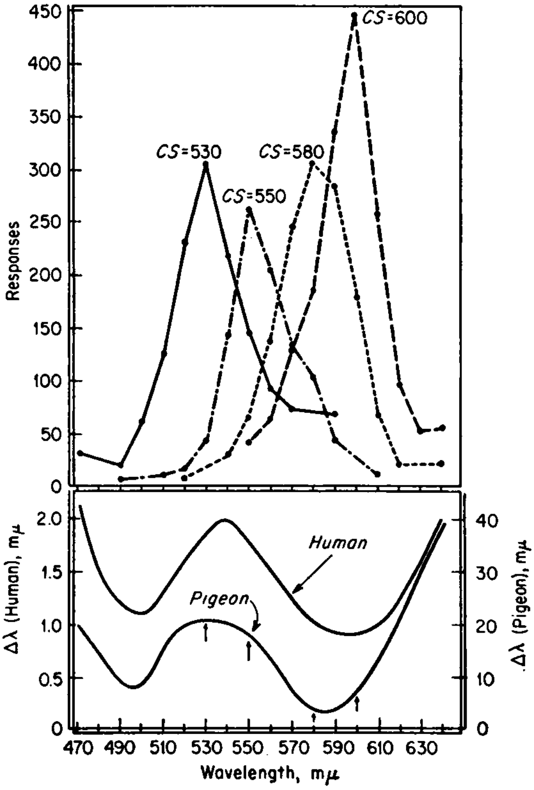

Fig. 1. Upper panel: Mean gradients from four groups of pigeons independently trained with the spectral values indicated at the top of each gradient. Lower panel: Hue discriminability functions obtained for humans, as adapted from Boring, Langfeld, and Weld (1948), and for pigeons according to Hamilton and Coleman (1933). The arrows on the latter function indicated the training values used by Guttman and Kalish. (From Guttman, 1963).

Guttman and Kalish (1956) had designed their experiment in order to address a specific theoretical issue, namely the relationship between discriminability and the slope of the generalization gradient. The major impact of their work, however, was to establish stimulus generalization as a productive area of research in its own right, particularly with the use of operant methods. Prior to their work, generalization gradients were respected more as theoretical entities than as empirical phenomena. Gradients served primarily as constructs for the explanation of other behavior, either as part of the apparatus required for general behavior theory (Hull, 1943), or for the explanation of particular phenomena such as conflict (Miller, 1944), choice behavior (Schlosberg & Solomon, 1943), and transposition (Spence, 1937). Various accounts of generalization itself were also proposed. These included the irradiation of excitation in the cerebral cortex (Pavlov, 1927), failure to discriminate test from training stimuli (Lashley & Wade, 1946), and response elicitation by elements common to both the training and testing situations (Bush & Mosteller, 1951; see also Mackintosh, 1974).

The theoretical analyses of stimulus generalization were premature at best, however, be-

---

The preparation of this manuscript was supported by Grant A0-102 from the National Sciences and Engineering Research Council of Canada and National Research Service Award No. 1-F32-MH08057 from the National Institute of Mental Health to W. K. Honig, and by National Research Service Awards Nos. 5-T32-EY07024 (National Eye Institute) and 1-F32-MH08018 (National Institute of Mental Health) to Peter J. Urcuioli. An earlier version of this review appears in German as "Reizgeneralisierung und Reizkontrolle" ("Stimulus Generalization and Stimulus Control") in H. Zeier (Ed.), *Pawlow und die Folgen*. Zurich: Kindler-Verlag, 1977. The authors gratefully acknowledge Janet Lord and Barbara Vavasour for their assistance in preparing the manuscript, and Donald S. Blough and Anthony A. Wright for their critical comments and stimulating discussions. Reprint requests may be addressed to either Werner K. Honig, Department of Psychology, Dalhousie University, Halifax, Nova Scotia, Canada B3H 4J1, or Peter J. Urcuioli, Department of Psychological Sciences, Purdue University, West Lafayette, Indiana 47907.

cause psychologists could not agree on the basic form of such gradients or even on whether so-called primary gradients existed in the absence of discrimination training (see the controversy between Lashley and Wade, 1946, and Hull, 1947). Early empirical work conducted within the framework of classical conditioning failed to have much impact because it provided neither reliable nor replicable data (Razran, 1949). Guttman and Kalish (1956) were aware of these conceptual and methodological problems, so their first approach to the study of generalization was the development of a procedure for obtaining reliable gradients from individual subjects. The procedure they adopted was one originally suggested by Skinner (1950):

> . . . experiments have shown that the pitch of an incidental tone, the shape of the pattern being [pecked], or the size of a pattern, if present during conditioning, will to some extent control the rate of responding during extinction. Some properties are more effective than others, and a quantitative evaluation is possible. By changing to several values of a stimulus in random order repeatedly during the extinction process, the gradient for stimulus generalization may be read directly in the rates of responding under each value. (p. 204)

Guttman and Kalish "read the gradients directly" from individual pigeons, and published a selection in their Figure 4 (p. 83). The gradients differed with respect to height, slope, and form, but their orderliness was most impressive: gradients could now be seen as empirical rather than theoretical entities.

One effect of Guttman and Kalish's pioneering efforts was to shift interest away from the underlying form of the generalization gradient to those independent variables which determined its characteristics. Research at the Duke laboratory, for instance, investigated stimulus generalization as a function of positive reinforcement at two or more spectral values (Kalish & Guttman, 1957, 1959), discrimination training between two spectral values (Hanson, 1959), extinction at a particular value (Honig, 1961), training along dimensions other than wavelength (Butter, 1963), and motivational level (Thomas & King, 1959).

This line of research was sufficiently influential that Guttman and Kalish (1958) published a popular description of their original work, and a chapter by Guttman (1963) addressing theoretical questions appeared in the series on *Psychology: a study of a science* edited by Koch. By the early 1960s, enough research effort had been devoted to stimulus generalization that David Mostofsky organized a conference on the topic at Boston University, the proceedings of which were later published (Mostofsky, 1965). A number of reviews treating various subtopics have subsequently appeared. Among them are chapters or articles by Terrace (1966a), Riley (1968, especially Chapter 2), Kalish (1969), Hearst, Besley, and Farthing (1970), Honig (1970), Thomas (1970), D. Blough (1975), Heinemann and Chase (1975), Rilling (1977), Mackintosh (1974, 1977), and Bitterman (1979).

It has now been 25 years since the original Guttman and Kalish work was completed. In view of its impact on the field of conditioning and learning, we think an appropriate tribute is an historical overview of the past quarter century of stimulus generalization research. Naturally, we cannot possibly summarize all of the empirical findings and theoretical considerations, so we concentrate instead on selected topics which emerged directly from the original Guttman and Kalish (1956) work, and which appeared to Guttman (1963) to be of particular theoretical significance. Our review falls into five sections: (1) Methodological and conceptual issues; (2) The genesis of stimulus control; (3) Discriminability and the slope of the generalization gradient; (4) Inhibitory stimulus control; (5) Attentional factors affecting the gradient. A common element in all of these topics will be the slope and shape of the generalization gradient, attributes which indicate the psychological properties of stimuli which control behavior.

METHODOLOGICAL AND CONCEPTUAL ISSUES

Stimulus Control and Stimulus Generalization

Stimulus control is defined in terms of specific results obtained under well-controlled test conditions. In the experimental analysis of behavior, a decremental generalization gradient is considered to be an index of stimulus control. This notion is incorporated in a familiar definition by Terrace (1966a):

> Stimulus control refers to the extent to which the value of an antecedent stimulus determines

> the probability of occurrence of a conditioned response. It is measured as a change in response probability that results from a change in stimulus value. The greater the change in response probability, the greater the degree of stimulus control with respect to the continuum being studied. (p. 271)

Several aspects of Terrace's definition are worth noting. First, the terminology is really appropriate to classical conditioning, which may reflect the fact that discrimination and generalization were initially studied by Pavlov and his associates. It is easy to reformulate the definition, however, with terms from instrumental learning, such as "instrumental behavior" and "response rate."

Second, the definition is clearly operational in character. As Mackintosh (1977) states:

> the term *stimulus control* has come to be used as a convenient shorthand expression describing . . . an observed relationship between changes in external stimuli and changes in recorded behavior. (p. 481)

This way of defining stimulus control implies a necessary comparison between two or more different conditions. The comparative procedure may be as simple as the presentation and removal of a stimulus in a discrete-trial conditioning procedure or as complex as an extended generalization test.

Third, it makes little sense to say that a particular response is controlled by a particular stimulus. The response may be elicited, suppressed, evoked, or directed by such a stimulus, but the concept of stimulus control, as defined here, refers only to a correlation between a set of stimulus values and a set of response values.

Fourth, since "stimulus control" serves a descriptive function, it does not have explanatory or theoretical properties. To say that a steep generalization gradient is the result of stimulus control is meaningless. Such a gradient is a measure of stimulus control. It is difficult in practice, however, to adhere to such restricted usage. Surely, some controlling process underlies orderly changes in response probability or response rate. Psychologists may not necessarily want to exclude theoretical significance from this concept. For example, one may wish to think of ways in which particular aspects of a stimulus are actually "processed" by the organism, such as the selection or perception of particular attributes of a complex stimulus (e.g., Reynolds, 1961b), or control by the associative values of a set of stimuli (Honig & Lindsay, 1975). These processes represent a contribution by the subject that is not encompassed by the neutral concept of stimulus control.

Guttman (1963) recognized this last point when he distinguished between the "physical stimulus" and the "functional stimulus." The former is described in physical terms as a value of the independent variable that defines the test dimension. The latter reflects the additional contribution of prior and current histories of reinforcement and the subject's sensory capacities (cf. Wright & Cumming, 1971). We hope that the meaning of this term—the functional stimulus—will emerge more clearly in the course of the present review. For the moment, consider that a stimulus is established as positive or negative through differential association with reward. The stimulus may also compete with other sources of behavioral control in the experimental situation. Furthermore, its discriminability with respect to other stimuli on the test dimension may not have any simple correspondence with the physical scale on which it is defined. All of these matters will be touched upon in this article. The critical point here is that the functional nature of the stimulus is discovered through the sort of control that it exerts, and that this control is readily and conveniently assessed with the procedure that yields a generalization gradient.

Area, Height, Slope, and Form

Generalization gradients can be described and compared on the basis of four characteristics: area, height, slope, and form. Comparisons are meaningful, of course, only if gradients are obtained under similar test conditions and assessed with the same measure.

The *area* is defined by the total gradient, which takes into consideration both the distribution of responses and the range of test values. *Height,* which is often confused with area, refers simply to the maximum level of responding along the gradient. Normally, this occurs at the training value, although it may vary depending upon the particulars of the training procedure. One should not assume that height and area share any necessary relation, because the latter is affected by gradient *slope*. Consider, for example, the gradients of Guttman and Kalish (1956) shown in Figure 1.

Clearly, the 600 nm gradient has the greatest height, but not the greatest area. This is because it is also the steepest of the four gradients. Slope measures the rate of change in response rate or response probability between two points along the gradient, one of which is usually the training value. Gradient slopes are not normally calculated, however, because the main concern in most studies is qualitative differences in slope rather than precise quantitative measures of individual gradients. When two or more gradients are compared, slope differences can be assessed statistically by the interaction term in analysis of variance. A simpler measure, although one more subject to error, is the percentage of responses emitted to the training stimulus; the greater this value, the steeper the gradient slope. In spite of the usual lack of quantitative specification, slope is considered to be the most sensitive index of stimulus control. A flat or horizontal gradient (one with zero slope) indicates the absence of stimulus control by the test dimension. Conversely, a decremental gradient (or an incremental gradient) shows that the stimulus dimension varied during testing controls to some degree the rate or probability of response.

Since gradients often differ in area as well as in slope, the question of appropriate comparisons becomes important. In such cases, slope is not an unequivocal index of stimulus control. For example, let gradient *T* incorporate all of the responses obtained during a test. Now delete half of the test periods at random to generate gradient *A*. The area of this gradient will be half that of gradient *T*, and its peak should be half as high. It will also be flatter. But does this mean that gradient *A* represents a reduced level of stimulus control? Probably not, since the flattening was due simply to a reduction in the amount of behavior sampled by the gradient. Similar considerations hold for more realistic conditions: two subjects, or groups of subjects, may for a variety of reasons respond at very different rates and thus produce gradients that differ in area and in slope. One solution to the problem of how to compare such gradients is to normalize them by expressing response output at each test value as a proportion (or percentage) of total test responses. Likewise, response output can be expressed as proportions of total responses given at the training stimulus, or *S*+. (The latter is probably less satisfactory, because it is more sensitive to fluctuations in total responding to *S*+.) When such a transformation is performed, gradients differing in area may well turn out to have the same slope; they are then said to be multiplicatively related.

Normalization of gradients by conversion to percentage measures is not, however, without its problems nor its critics (Morgan, 1969). It seems a reasonable procedure when gradients are multiplicatively related, because one can then argue that stimulus control is basically the same for all gradients. However, consider the case where gradients are additively related: they have the same actual slope, but they differ by some constant amount at each value. Such gradients are parallel, but if they are normalized, the higher gradient necessarily becomes flatter. Is it then justified to conclude that the higher gradient reflects less stimulus control? The answer is not obvious. This matter will be taken up again in our discussion of attentional factors, where it has led to some dispute.

Area, height, and slope are rather clearly defined, but the concept of *form* is less precise, if not less important. The form of the gradient depends largely on the training procedure carried out in advance of testing or, in some cases, in the course of testing. It also depends on the psychological characteristics of the test stimuli—their discriminability, their similarity to the training value, and so forth. Much of the rest of this paper will be devoted to a discussion of the determinants of gradient form, so a couple of examples will suffice here.

If the stimuli lying to either side of *S*+ on the test continuum are not equally similar to *S*+, then the gradient will be *asymmetrical*, being flatter on the side where the stimuli are more similar to *S*+. If subjects are trained to discriminate between two values on a single dimension (e.g., one value serving as *S*+, the other as *S*—), then the gradient will probably show a ***peak shift:*** the greatest number of responses will not occur at *S*+, but at a test value more distant from *S*— than *S*+. Finally, if the stimulus of interest signals a period of extinction (*S*—) and lies on a continuum orthogonal to *S*+, then the gradient around *S*— will appear inverted or *incremental*. These examples will be illustrated later in our discussion of the training and testing conditions which generate them.

Procedural Variations in Testing

Guttman and Kalish (1956) presented their test stimuli in randomized order for 30-sec trials in extinction. All birds in a group were tested with each stimulus value. This has been the most common test procedure for obtaining generalization gradients, although both practical and theoretical considerations have led to certain variations, some of which are discussed below.

Single-stimulus vs. multiple-stimulus tests. As indicated above, one question that was in dispute prior to much of the empirical work on generalization was whether the gradient is fundamental, or whether it emerges because the subject has a chance to compare test stimuli (Hull, 1947; Lashley & Wade, 1946). One way to distinguish between these alternatives is to eliminate comparisons among the training and test stimuli. This can be accomplished by testing subjects at only one value, a procedure which is admittedly expensive and cumbersome.

Two notable studies, however, used the single-stimulus test procedure to obtain gradients for comparison with those obtained from the usual within-subjects, multiple-stimulus procedure. Hiss and Thomas (1963) trained four groups of pigeons to respond to 550 nm and then tested them in extinction with either 550, 530, or 510 nm. The comparison group was tested with all three values. Both the single-stimulus and comparison gradients were orderly, although the latter was somewhat steeper when plotted in terms of mean response rates. Kalish and Haber (1963) also trained pigeons to respond to 550 nm, but tested with six rather than three values, ranging from 550 nm to 490 nm. They also obtained orderly response gradients when the single-stimulus test data were combined. For comparison purposes, the authors used the left half of the 550-nm gradient reported by Guttman and Kalish (1956; see Figure 1). The Guttman and Kalish gradient was steeper than their single-stimulus gradient, although for various reasons the former was probably inadequate for comparison. In any case, both their study and that of Hiss and Thomas (1963) clearly demonstrate that, at least with pigeons and spectral stimuli, decremental generalization gradients can be obtained without sequential, juxtaposed presentations of different test values.

Resistance to reinforcement. Certain training procedures produce very low rates of responding, particularly if the stimulus of interest is correlated with extinction. Response output during testing may be so low as to mask any differential effects along a dimension. To correct this sort of problem, Hearst, Besley, and Farthing (1970) proposed a "resistance to reinforcement" procedure, in which a standard reinforcement schedule is in effect during presentation of all test stimuli. Dimensional control is revealed if responding recovers at different rates with different test values. For example, incremental gradients are produced if *S*−, and those stimuli most similar to it, occasion slower rates of recovery than more remote values. This test can be particularly helpful in situations of low baseline response rates, not only because of its diagnostic value, but also because repeated testing can be conducted without the threat of extinction.

Steady-state testing. Generalization gradients obtained during extinction can be viewed as representing differential resistance to extinction at different test values. Extinction necessarily limits the amount of test data that can be collected, however, and produces its own changes in the slope and form of the gradient (Friedman & Guttman, 1965).

One way to counteract undesirable extinction effects (and to prolong the length of time for data collection) is to continue to reinforce responding to the training value(s) during testing, albeit intermittently (steady-state testing). This amounts to a "discrimination test" rather than a generalization test and, of course, it will also affect the slope and form of the gradient, as responding to all but the reinforced value extinguishes. However, the transitional effects occur more slowly than with testing in extinction. Furthermore, if the object of study is generalization to a set of closely spaced stimuli, the discrimination problem is pretty much overcome (e.g., D. Blough, 1969a, 1972, 1975, 1978; P. Blough, 1972).

A particularly interesting form of the steady-state procedure was developed by D. Blough (1975). For reasons which are not of concern here, Blough established a baseline level of responding to 25 values within a restricted spectral range. Individual trials lasted for 20 sec, after which time the pigeon could procure a conditioned reinforcer. The conditioned reinforcer was periodically followed by food, assigned at random to all stimuli within the test range. Once the baseline was established, one or more stimuli were selected for special treatment. For example, Blough provided additional reinforced trials at the central value of the stimulus set in order to generate a positive (decremental) gradient. Conversely, additional unreinforced trials at the same value were used to generate a negative (incremental) gradient.

This steady-state procedure also revealed interesting changes in the form of the gradients across successive quarters of the 20-sec trial period. A typical positive gradient, broken down into its 5-sec composites, is seen in Figure 2. Responding during the first two quarters of a trial yields steep orderly gradients. The gradient becomes higher and flatter, however, toward the end of a trial, when reward is imminent and response rates are high. This variation in gradient form suggests that stimuli other than those on the key (perhaps prior key pecking and collateral behavior) come to control responding in free-operant situations (cf. D. Blough, 1963). The effects of different reinforcement schedules upon dimensional stimulus control provide some additional evidence in this regard.

Schedule Effects

Most operant research on stimulus generalization has involved VI schedules in training and exposures for 30 sec or more to each stimulus in testing. Training schedules other than VI, however, are known to produce special effects on generalization gradients. Thomas and Switalski (1966), for example, reinforced pigeons for responding on a variable-ratio (VR) schedule prior to generalization testing. The VR schedule generated very high rates of responding. Pigeons in a yoked-control group could obtain reinforcement whenever reward was provided for the VR birds. This arrangement matched the frequency and distribution of rewards in the VR group but did not involve a contingency based on the number or rate of responses. Thomas and Switalski found that spectral gradients obtained from the VR birds were considerably flatter than those of the yoked birds. One interpretation of this effect is that the cues associated with high rates of responding masked or overshadowed control by color on the key. (A fuller discussion of masking and overshadowing is provided later in this review.)

Hearst, Koresko, and Poppen (1964) obtained gradients on the dimension of line orientation after training pigeons on a differential-reinforcement-of-low-rates-of-responding (DRL) 6-sec schedule, which provided reward only following interresponse times (IRTs) greater than 6 sec. The gradients were very flat compared to those obtained with training on standard VI schedules. This difference was not due to a low rate of reinforcement in the DRL condition. Birds who obtained a similar frequency of reward on VI schedules yielded good decremental gradients. The authors suggested that exteroceptive stimulus control was reduced in the DRL group by frequent pausing and the occurrence of stereotyped chains of collateral behavior engendered by the training schedule. A more recent study by Gray (1976), however, offers a different interpretation.

In Gray's study, key pecking to a 570-nm stimulus was reinforced on a DRL 8-sec schedule. After extended training, each pigeon received steady-state generalization tests with closely spaced spectral values. The two birds that performed poorly on the training schedule provided steep decremental gradients. The

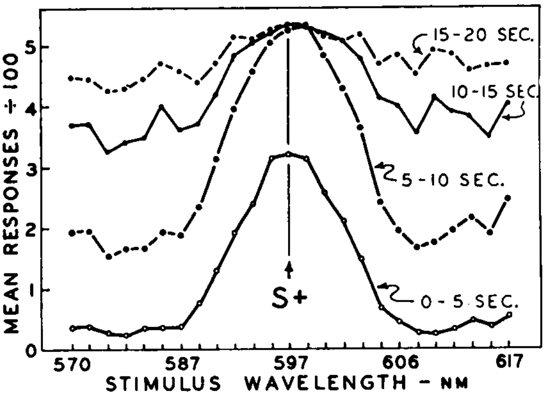

Fig. 2. Spectral generalization gradients from a single pigeon obtained across successive quarters of 20-sec trials. Gradients are averaged across four sessions in which the probability of reinforcement at all test stimuli was .167 and additional reinforced trials were presented at 597 nm. The gradients become higher and flatter over the course of the fixed interval. (From D. Blough, 1975).

three birds that showed better control by the schedule contingencies, on the other hand, provided much flatter gradients. Gray then separated the test responses into "bins" defined by the IRT preceding each response, and plotted "mini-gradients" based upon this analysis. The birds that provided the steep gradients (and performed poorly on the DRL schedule) showed decremental generalization in all of the mini-gradients. In contrast, those birds whose responding indicated strong control by prior IRTs provided decremental gradients only for responses preceded by IRTs that met the criterion for reinforcement. Their other mini-gradients were either flat or inverted. This indicates in effect a conditional discrimination: responding was under the control of the stimulus on the key only in the presence of cues arising from IRTs that could be reinforced. Thus, long pauses and collateral behavior that accompany DRL schedules may not actually reduce control by the discriminative stimulus. On the contrary, the flat gradients obtained by Hearst, Koresko, and Poppen (1964) could have reflected the additional contribution of responses with IRTs that did not meet the requirements for reward and would not therefore be associated with that event.

Gray's experiment points to the importance of the correlation between responding and reinforcement in the development of stimulus control. Further evidence in this regard is provided by a second study by Hearst, Koresko, and Poppen (1964). They trained different groups of pigeons on VI schedules with reinforcement parameters ranging from 15 to 120 per hour. The slopes of their corresponding gradients of line orientation were directly related to reinforcement density during training. More specifically, richer schedules of reinforcement created steeper gradients of generalization. Although response rates differed substantially between groups during training, total response output during testing did not. Consequently, the slope differences could not be attributed simply to differences in response rates. Presumably, the rate of reinforcement in conjunction with the instrumental response determines the degree to which the training value acquires control of responding. It is no accident that in most studies of stimulus generalization, the favored training schedules have ranged from VI 30-sec to VI 120-sec.

Other Dependent Measures

Response rate and response probability are the traditional dependent variables in the study of generalization. They have been popular no doubt because they generate orderly, reproducible results with respect to a variety of independent manipulations. Nonetheless, other measures can provide orderly gradients and even additional information unavailable within traditional paradigms.

One promising alternate measure is stimulus duration, the amount of time that a subject spends in the presence of each test stimulus. Naturally, the subject must have control over the length of time each stimulus is on, so a control key (in addition to or instead of the usual food key) is provided. A response to this key terminates the current stimulus and initiates the next in the sequence. It is often useful in this situation to limit the duration of any particular trial and to advance the test sequence automatically when this limit is reached. Responding to the food key (if available) can be recorded as well.

Beale and Winton (1970) used the stimulus duration measure to demonstrate incremental gradients around a stimulus value associated with extinction. In fact, their duration gradients were steeper than gradients of local response rate obtained in the same experiment. Honig and Beale (1976) extended their work to demonstrate decremental gradients of stimulus duration around *S*+ and showed that such gradients can be predicted from rates of responding independently obtained with the same test stimuli.

Three points concerning stimulus duration should be noted. First, this measure is particularly helpful in assessing dimensional control by stimuli associated with extinction, because it circumvents the problems associated with low response rates (i.e., floor effects). Second, stimulus duration provides a convenient index of stimulus control in situations where food-reinforced responding is absent (e.g., if reward is presented independently of responding), or if the discriminative stimuli are diffuse (such as tones). Finally, discrimination learning is facilitated when subjects have control over the discriminative stimuli. These considerations suggest that stimulus duration may in fact be a more useful measure of stimulus control than response rate [cf. Honig, Beale, Seraganian, Lander, & Muir (1972)].

Stimulus Generalization as a Psychophysical Problem

A recent approach to the analysis of stimulus generalization comes from the application of statistical decision theory (D. Blough, 1967, 1969a; Heinemann, Avin, Sullivan, & Chase, 1969; Heinemann & Chase, 1975). Essentially, generalization is viewed as a psychophysical problem in which the subject's task is to detect and to report changes in the stimulus previously correlated with reinforcement. Performance is thus considered to be a joint function of the subject's sensitivity to variation along the test dimension and its criterion or threshold for responding. The advantage of this conceptualization is that changes in gradient form or slope can be interpreted either as a result of changing sensitivity, changing response criteria, or both. Its usefulness is exemplified by the work of Wright (1972, 1974) on hue discrimination and spectral generalization in the pigeon, a topic we will discuss more fully later in this review. The theory is best applied to discrete-trial situations with response probability as the dependent variable, but it is nonetheless instructive to consider how operant generalization might be construed from a signal-detection perspective.

Statistical decision theory models the joint effects of sensitivity and response criteria in the following manner. The training value and all other values along the test dimension give rise to discriminal processes (neural or stimulus effects), which vary somewhat from trial to trial. This variation is usually assumed to take the form of a normal distribution, although this is neither a necessary assumption nor is it critical to detection and decision processes. Each test stimulus (including the training value) produces its own distribution, which is displaced from the others along the psychological axis of stimulus effects. During training, the subject establishes criteria for responding (cf. Boneau & Cole, 1967) such that any stimulus effect which falls within the criterial zone generates the pattern of responding associated with reinforcement. Similarly, during the generalization test, responding to each test stimulus is determined by the probability that its associated effect falls within this response zone. Test stimuli close in value to the training stimulus produce a relatively high probability of response on any given trial because a greater proportion of the area under their stimulus-effects distributions falls within the response zone of the training distribution. Conversely, test values remote from the training value produce a relatively low probability of response because there is less areal overlap. Sensitivity to differences among stimuli (discriminability) can thus be seen to affect gradient slope. We should emphasize that sensitivity refers to the scale of functional rather than physical stimuli, because the former must take into account the necessary sensory transduction of the latter.

The slope of the generalization gradient is also affected by how lax or how strict the subject's response criteria are. Lax criteria will generally produce broad or shallow gradients because a greater proportion of the area under each test distribution then falls within the response zone. Strict criteria, on the other hand, make the gradient sharper or steeper because the response zone cuts off relatively smaller areal proportions. The modulation in gradient slope caused by variation in response criteria mimics in some respects the effects caused by variation in the subject's sensitivity. Certain changes in the gradient, however, are more reasonably modelled by criteria shifts than by sensitivity differences. For example, the finding that generalization gradients become lower in height and steeper in slope during prolonged testing in extinction (Friedman & Guttman, 1965; Hoffman & Fleshler, 1961) can best be accounted for by criteria which grow increasingly strict. A similar argument might explain the fact that gradients along one dimension become steeper following discrimination training along an orthogonal dimension (Honig, 1969, 1974; Turner & Mackintosh, 1972). D. Blough (1969a) has extended this sort of analysis to show how gradient summation following training with two reinforced wavelengths can be predicted. Finally, Heinemann and his colleagues argue convincingly that criterion changes can explain variation in gradient slope due to different degrees of discrimination difficulty during training (Heinemann et al., 1969) and to the presence vs. absence of redundant stimuli which signal the prevailing contingencies of reinforcement (Chase & Heinemann, 1972; see also D. Blough, 1972).

THE GENESIS OF STIMULUS CONTROL

One of the striking features of the large body of data on stimulus generalization is the orderliness of the gradients—that is, the monotonic relationship between decrement in responding and the difference between training and test values. Guttman paid tribute to this characteristic in his review of 1963:

> The curve of generalization is not just an "error curve" (in the statistical sense) but a curve of errors systematically regulated and distributed in a nice relation to physical facts. Everything the subject does except respond to the conditioned [stimulus] is an error, but the errors have a pattern and sense, and they reveal a structure. The animal has a "color-space" which is in part a map of the measurable (by us) attributes of the various stimulations. This color space . . . is the structure, the system, the organization within which the stimulus class is determined. (p. 144f)

These remarks, based upon spectral generalization gradients, could now be extended to many dimensions, including line orientation, auditory frequency, floor tilt, and rate of alternation, to name a few. Even when stimulus control on a particular dimension has to be established through discrimination training, responding to novel test values indicates a sensory order. Nonetheless, the question arises to what extent this sensory order depends upon prior experience with the relevant stimulus dimension.

This question was debated by Lashley and Wade (1946) and Hull (1947) some years before Guttman and Kalish carried out their research. Hull thought that at least some gradients were primary, independent of prior discrimination training on the test dimension. Such gradients formed an integral part of his theory of behavior. Lashley and Wade maintained the opposite view, namely that decremental generalization gradients depended upon prior differential experience with values from the test dimension. Guttman's (1963) view was closer to Hull's:

> If we raise the ontogenetic questions, I do not believe that the pigeon "learns" the spectral order, or any sensory order. This is a sheer guess, but for the rhesus monkey, Ganz has determined that orderly spectrum generalization exists in animals raised in darkness from birth. Surely there must be some stimulational structures biologically given which can be built upon and combined with others. (p. 146)

Guttman's reference to Ganz involves one of two influential studies on the effects of restricted early rearing on stimulus generalization. Such methods are required to obviate differential reinforcement on the test dimension before generalization gradients are obtained. Ganz and Riesen (1962) reared two groups of macaque monkeys, one in darkness and the other in normal illumination. All animals then learned to press a lever for reinforcement while one eye was illuminated with diffuse monochromatic light through a contact lens. During seven daily tests, reinforced presentations of the training value were mixed with test stimuli presented in extinction. In general, gradients from dark-reared subjects were steeper, indicating that prior differential experience had *reduced* the control by the spectral stimuli. Unfortunately, the testing procedure itself provided an opportunity for discrimination learning, so this effect may have been artifactual. In fact, during the first test session, gradients from both the normal and the dark-reared groups were quite flat.

Peterson (1962) carried out a much cited study with ducklings which led to the opposite conclusion regarding the genesis of sensory order. He raised four subjects in an environment illuminated only by a sodium vapor lamp, which emits radiation of 589 nm in the visible spectrum. Two control subjects were raised in white light. Peterson then trained all ducklings to respond to a key illuminated with a 589 nm light. (The reinforcer was water.) A generalization test on the spectral continuum revealed a flat gradient for the monochromatically reared group and decremental gradients for the control animals.

Although Peterson's method seemed promising for further study of the genesis of stimulus control, replications of his experiment have been unsuccessful (Malott, 1968; Mountjoy & Malott, 1968; Tracy, 1970). Rudolph, Honig, and Gerry (1969), for example, raised quail and chickens in conditions which included restricted portions of the spectrum, monochromatic illumination, and darkness, and found decremental gradients that were generally steeper than those of white-reared controls. Their last experiment was a systematic replication of Peterson (1962). Two groups of chickens were reared either in the presence of a sodium vapor lamp or in white light, with other aspects of the environment carefully controlled. Following key-peck training to 590 nm, tests were carried out at 20-nm intervals between 510 and 590 nm. Mean gradients are presented in Figure 3. (The data for one sodium-reared animal who emitted less than 100 responses on the test have been omitted, since such low response rates artifactually produce steep gradients.) Clearly, the mean gradients are almost identical; if anything, the one from the sodium-reared group is a bit steeper. The reasons why Peterson (1962) obtained the opposite result are unclear. His data may reflect a difference in the level of overall illumination, temperature, or contrast in the home cage where incandescent lamps presumably produced more light and heat than the sodium vapor lamps.

These findings were extended to the effects of explicit discrimination training in two independent but very similar studies by Rudolph and Honig (1972) and Terrace (1975). The former raised chicks and the latter raised ducklings in monochromatic light of 589 nm. Subjects in both experiments were then trained on a successive discrimination with 589 (or 590) nm as $S+$. $S-$ was 570 nm in the Rudolph and Honig experiment, whereas Terrace used two negative values, 570 and 610 nm, with different groups. Both studies found that monochromatic rearing did not affect discrimination learning, relative to controls raised in white light. Furthermore, all of the groups (and most of the individual subjects) showed peak-shifted generalization gradients, with maximum responding to a test value displaced away from $S+$ in a direction away from $S-$.

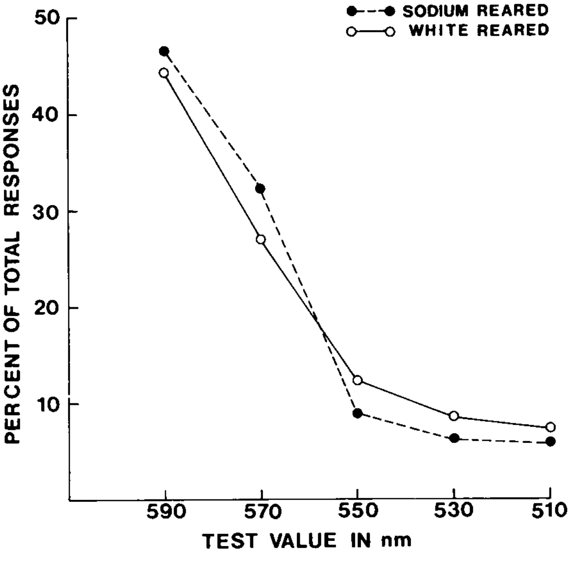

Fig. 3. Mean gradients obtained from chicks reared either in monochromatic light from a sodium vapor lamp or in white light. (Based on data from Rudolph, Honig, and Gerry, 1969).

Taken together, the early rearing data suggest that sensory order is primarily dependent on physical maturation rather than differential experience. This suggests that the "stimulational structures biologically given" should be discoverable. Regrettably, little progress has been made in establishing connections between generalization gradients and neural structure or function, in spite of the enormous progress in sensory physiology over the last 25 years. A topological correspondence between cortical and skin surface may be evident for tactile stimuli, but the correspondence in other modalities may be quite different (Zeki, 1980). We know that individual neurons in the visual cortices of many species respond best to a preferred orientation of a line or direction of movement, and there may be a rough gradient of responsiveness around the preferred value. But these patterns are not the outcome of the acquisition of stimulus control, except in the general sense that early experience seems to determine the statistical distributions of cells which favor particular orientations. A true correspondence between gradients acquired through training and neural responsiveness to the test stimuli still awaits demonstration.

DISCRIMINABILITY AND THE SLOPE OF THE GENERALIZATION GRADIENT

The original purpose of Guttman and Kalish's (1956) experiment was to study how stimulus generalization relates to stimulus discrimination. Guttman and Kalish suggested, as had others (e.g., Lashley & Wade, 1946), that these two processes should be the inverse of one another: that is, subjects will generalize to the extent that they cannot discriminate. This somewhat intuitive notion (later called the *inverse hypothesis*) predicts that the slope of the generalization gradient between any two points along a stimulus continuum will be inversely proportional to the size of the corresponding difference threshold. In other words, gradient slopes should be shallow where the difference threshold is large (stimuli are relatively similar) and steep where the threshold is small (stimuli are relatively dissimilar).

Guttman and Kalish (1956) tested these and related predictions by selecting a stimulus dimension for which the difference threshold (discriminability) function had already been established—namely, wavelength. Hamilton and Coleman (1933) previously determined a hue discriminability function for pigeons (see lower panel of Figure 1), one which showed considerable variation in the size of the differential threshold across the spectral range from 470 to 630 nm. Consequently, Guttman and Kalish chose their four training values (530, 550, 580, and 600 nm) in order to reveal local changes in hue discriminability across the spectrum. The 530-nm stimulus, for example, was selected from a region where thresholds were relatively large, so the gradient around that value was predicted to be relatively shallow. Training with 580 nm, on the other hand, was expected to produce a relatively steep gradient, since it was located in a spectral region of small difference thresholds. Finally, the gradients around 550 and 600 nm were predicted to be asymmetrical in slope, since thresholds toward either side of these values changed in opposite directions.

The empirical gradients obtained by Guttman and Kalish (upper panel of Figure 1) did not, however, confirm these predictions. For the most part, the slopes of the average gradients were similar across the four training conditions. Differences that were observed ran contrary to those predicted. The gradient around 580 nm, for instance, was no more sharply sloped than the one around 530 nm. The steepest gradient occurred between 600 nm and the longer wavelengths where discriminability was supposedly very poor (i.e., thresholds were quite large). The 550-nm gradient was clearly asymmetrical but in the wrong direction. Kalish (1958) later summarized these discrepancies by comparing threshold values estimated from the generalization data with those read directly from the pigeon's hue discriminability function. The former were derived by intersecting each of the four gradients at different response levels, recording the band width (in nm) at each level, and then averaging across levels. These values were then examined against the four corresponding difference thresholds reported by Hamilton and Coleman (1933). The results of this analysis can be seen in Figure 4. Kalish's derived values are shown as the filled circles, Hamilton and Coleman's threshold data as the unfilled triangles. (The remaining set of data points will be discussed later.) The discrepancy between sets of data is obvious, with marked departures from coincidence following training at the two longest wavelengths.

Methodological Problems in the Test of the Inverse Hypothesis

Guttman and Kalish's (1956) generalization results clearly questioned the validity of the inverse hypothesis. The authors could have reasonably rejected the hypothesis, but they were careful to point out that there were at least two factors which could have simply masked the inverse relationship: (1) brightness variation among their stimuli, and (2) procedural differences between their generalization technique and typical discrimination techniques.

The first, brightness variation, was an important evaluative concern since the wavelength stimuli used by Guttman and Kalish (1956) were not equated for brightness. Consequently, if this variable exerted any degree of control over responding, then the test data did not provide an unbiased assessment of per-

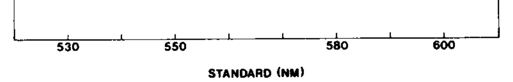

Fig. 4. Three measures of the pigeon's sensitivity for spectral differences at the training values used by Guttman and Kalish (1956): gradient band width, minimum difference thresholds (Hamilton and Coleman, 1933), and psychophysical difference thresholds at $d' = 2.0$ (Wright, 1972). The scale on the right-hand ordinate is applicable to Wright's data only. See text and Figure 5 for further explanation.

ceptual similarity among wavelengths. The local steepness of the 600-nm gradient toward the longer wavelengths is a good illustration of this point. Its sharp slope might be explained by decreasing brightness in that region of the spectrum. D. Blough (1961), however, ruled out brightness variation as an explanation for Guttman and Kalish's (1956) results by replicating a good deal of their work using monochromatic stimuli corrected for both the sensitivity of the pigeon's eye (D. Blough, 1957) and the output of the monochromator. Blough found that the spectral regions where pigeons generalized the most did not correspond to areas where discrimination was supposedly poorest, and vice versa.

The second possibility entertained by Guttman and Kalish (1956) was that procedural differences between generalization and discrimination tests produced the observed discrepancies. The former are typically conducted after training to a single spectral value, whereas the latter involve differential reinforcement to simultaneously or successively displayed values. Comparative judgments between stimulus pairs, therefore, are not required for generalization, whereas they are an essential ingredient for discrimination. This difference alone may preclude detection of an inverse generalization-discrimination relationship, as some later research tended to suggest.

Kalish (1958), for example, showed human subjects one of four spectral values (training) and then presented a series of novel test stimuli (generalization). Subjects were told to respond to the test stimuli only if the individual values were identical to the one training value. Kalish reported that the forms of the resulting generalization gradients corresponded closely to what would be predicted on the basis of human hue discriminability data.

Marsh (1967) investigated procedural effects with pigeons by systematically varying the number of stimuli presented during generalization tests. Marsh trained separate groups to respond to 520 and 590 nm and then tested in extinction with either two, three, or four spectral values (including the training value). With two test stimuli, the slopes of the 520- and 590-nm gradients were in the same proportion as their corresponding difference thresholds. The proportionality broke down, however, with greater numbers of test stimuli, primarily because the 520-nm gradient steepened. Marsh's two-stimulus generalization test condition was, of course, similar to discrimination training, the only difference being the absence of differential reinforcement. The three- and four-stimulus tests, on the other hand, differed both in the number of discriminative stimuli and in the contingencies of reinforcement.

Ganz (1962) tested the inverse hypothesis with rhesus monkeys using a complete within-subjects experimental design and discrimination training (as opposed to single-stimulus training) prior to generalization testing. Ganz taught his monkeys a series of three successive discriminations between pairs of stimuli drawn from the spectral set of 449, 509, 567, and 631 nm. Adjacent values within the set served as *S*+ and *S*− for each discrimination. (Absolute stimulus values did not change across successive training phases.) Following each discrimination, the monkeys received a generalization test in extinction with test stimuli consisting of the current positive and negative training values and five novel values lying between them. Ganz found that local differences in the post-discrimination generalization gradients reflected local changes in the monkey hue discriminability function, For example, gradient slopes were relatively steep near *S*+'s located in spectral regions of high discriminability (e.g., 567 nm) and were relatively shallow near *S*+'s located in low discriminability regions (e.g., 449 nm).

The aforementioned studies suggested one possible reason for Guttman and Kalish's failure to confirm the inverse hypothesis. We should note, however, that two of these studies (Ganz, 1962; Kalish, 1958) were done with other species, so their applicability to the Guttman and Kalish (1956) research must be assumed. Marsh's (1967) study was more comparable to that of Guttman and Kalish, although it was not nearly as extensive. Regardless of final conclusions, there is at least one other major methodological factor which may have been responsible for Guttman and Kalish's (1956) negative findings, although not discussed by them. This concerns the measure of hue similarity itself—namely, rate of responding. As we mentioned previously, response rate is affected not only by the characteristics of the test stimuli but also by variables more difficult to bring under experimental control. For example, responses following short IRTs (400 msec or less) are controlled less strongly by exteroceptive stimuli than those following longer IRTs (D. Blough, 1963, 1969a). Thus, changes in rate which might reflect perceptual differences between test stimuli may be overshadowed in part by the cues controlling short-IRT responses (see also D. Blough, 1965). In particular, where response rates are high and short IRTs predominate, a greater proportion of responses are not directly controlled by the stimuli of interest. The problem is largely eliminated at lower rates since the relative frequency of longer IRTs typically increases. Nonetheless, slope differences are usually assessed near the training value where response rates are usually high. Discrete-trial procedures can be used to avoid this problem, although cues associated with the beginning of a discrete trial may then compete for behavioral control (cf. Baker & Holland, 1968; D. Blough, 1969a, 1978).

The Nature of the Psychophysical Function

Up to this point, we have focussed on potential problems associated with Guttman and Kalish's (1956) generalization paradigm. We have not yet considered the discriminability findings reported by Hamilton and Coleman (1933). If their data were inaccurate, then Guttman and Kalish's failure to confirm the inverse hypothesis would not be surprising and most of what has already been discussed will be moot. In fact, the Hamilton and Coleman data *are* inaccurate, although this was not firmly established until the early 1970s, 15 years after Guttman and Kalish's pioneering efforts.

Hamilton and Coleman (1933) obtained their differential threshold function by training pigeons to choose between two patches of spectral light in a Lashley-type jumping stand. The threshold at each particular wavelength was defined as the minimum spectral difference which maintained accurate discrimination performance. Hamilton and Coleman found that absolute thresholds measured in this fashion were larger for pigeons than corresponding thresholds for humans. The relative functions, however, were similar for both species (see Figure 1), a finding which apparently promoted acceptance of their data (cf. Kalish, 1958, p. 643).

Suspicions about the discriminability function began to arise, however, when others found consistencies among the forms of certain generalization gradients. D. Blough (1961), for instance, reported that his 550-nm gradients were asymmetrical with relatively steep slopes toward shorter wavelengths, much like those obtained by Guttman and Kalish (1956). The 550-nm asymmetry had also been obtained by Hanson (1959, 1961), Honig, Thomas, and Guttman (1959), and Thomas and King (1959). It was later replicated by Butter (1963), Friedman (1963), Holland and Baker (1968), and Thomas and Switalski (1966). The 555-nm gradient reported by Tomie, Davitt, and Engberg (1976) is also similarly shaped. The hue discriminability data from Hamilton and Coleman (1933) predict the opposite asymmetry. D. Blough (1961) also found that his 570-nm gradient was rounded and symmetrical much like Guttman and Kalish's 580-nm gradient. This form also runs counter to the prediction of a relatively sharp gradient which, if asymmetrical, should be steeper toward longer wavelengths.

These and other discrepancies can be reconciled, however, with pigeon hue discrimination data recently obtained by Wright (1972, 1974), who applied a signal detection approach to this problem. Wright trained his birds in a yes/no conditional discrimination where they detected the presence vs. absence of a hue difference in a split field. One reference wavelength appeared in the left half of the field during a single experimental session, whereas six comparison wavelengths appeared on different trials in the right half of the field. The comparisons were either shorter in wavelength than the reference or equal to it. After the birds had been trained on the basic procedure, isosensitivity functions (Engen, 1971) were generated for each wavelength difference by varying the relative probability of reinforcement for correct same vs. correct different choices. A bias-free index of discriminability, d', was then computed from each function at the point of equal bias (i.e., at the point where the probabilities of correct same vs. correct different responses were equal). These d' values, plotted against their corresponding wavelength differences, yielded linear psychometric functions whose slopes reflected hue discriminability across the spectrum. Steep slopes indicated good discriminability, shallow slopes poor discriminability.

Sets of equally discriminable wavelength differences were then derived by intersecting the

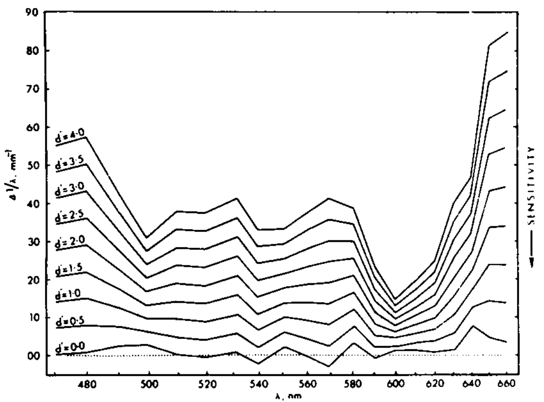

Fig. 5. Mean relative sensitivity for hue discrimination for four pigeons as indicated by equal *d'* contours. The wavenumber difference at each spectral value is that difference which, on the average, produced a level of discrimination performance corresponding to the plotted *d'*. (From Wright, 1974).

psychometric functions at various *d'* values. Figure 5 shows these data plotted against reference wavelength. The equal *d'* contours represent relative spectral sensitivity or hue discriminability for Wright's pigeons. The minima at 500 nm, 540 to 550 nm, and 600 nm represent spectral regions where hue discriminability is best. The areas between these values as well as the shorter (<500 nm) and longer (>600 nm) wavelength regions show somewhat poorer discriminability.

Wright's data together with the inverse hypothesis predict that generalization gradients should be relatively steep around 500 nm, 540-550 nm, and 600 nm, where differential thresholds are relatively small. Gradients should be relatively shallow between these values and in the extreme spectral regions, where thresholds are relatively large. If Guttman and Kalish's (1956) results are evaluated against these new predictions, then good evidence for an inverse generalization-discrimination relationship is obtained. This can be seen in Figure 4 where Kalish's (1958) discriminability estimates based on Guttman and Kalish's generalization curves (filled circles) are shown together with the corresponding values read directly from Wright's $d'=2.0$ contour (unfilled squares). The agreement between the relative functions is surprisingly good considering that only four pairs of points are being compared, and that Guttman and Kalish's training values were selected on the basis of other data (i.e., Hamilton & Coleman, 1933).

A more formal test of the predictions generated by the new discriminability findings was provided by P. Blough (1972). In one of her experiments, separate groups of pigeons were initially trained with one of six different spectral values selected from regions with varying discriminability characteristics. The birds were subsequently tested in a steady-state procedure. Two test series were run, one with stimuli spaced 4 nm apart (earlier data), the other with stimuli spaced 2 nm apart (later data). Figure 6 shows that the forms of the resulting generalization gradients differed in a way consistent with the inverse hypothesis. For example, both the 540- and 600-nm gradients are steep and symmetrical, with particularly sharp slopes in the latter condition. The 570-nm gradients are also symmetrical but relatively shallow. The gradients around 630 nm are steeper toward 600 nm, whereas those around 510 nm are steeper toward 540 nm. The only anomalies come from the test data obtained in the vicinity of 500 nm, where gradients should be steep but are actually quite flat.

Jitsumori (1978) recently provided another test of the inverse hypothesis using training and testing procedures similar to those used by Ganz (1962). Separate groups of pigeons learned successive discriminations between 420 and 470, 470 and 520, 520 and 570, or 570 and 650 nm. A yes/no choice procedure was used in conjunction with the successively presented

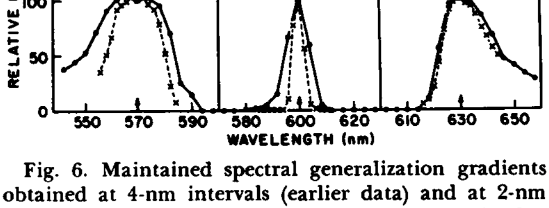

Fig. 6. Maintained spectral generalization gradients obtained at 4-nm intervals (earlier data) and at 2-nm intervals (later data) from pigeons. The reinforced value for each training condition is indicated by the arrow. Differences in slope and symmetry depend upon the location of the training value along the spectral continuum. (From P. Blough, 1972).

spectral stimuli. Following discrimination training, pigeons received a series of maintained discrimination tests where presentations of a single novel stimulus were intermixed with presentations of the training stimuli. Different novel stimuli were used in different tests, but all were drawn from the spectral regions bounded by the training values. Discriminability indices were derived from the test data by summing the differences between choice probabilities at adjacent wavelengths and computing corresponding reciprocals. A low reciprocal value indicated good hue discriminability in the spectral region from which it was derived; a high value indicated poor discriminability. Jitsumori reported that his indices were lowest around 450 nm, 500 to 510 nm, 530 to 540 nm, and 600 nm. The latter three values correspond closely to the minima in Wright's discriminability function.

Finally, it is of some interest to note that Shepard (1965) had earlier suggested that hue discrimination for pigeons was best in the spectral regions where Wright (1972) later observed minima in the discriminability function. By a unique transformation of the wavelength scale, Shepard was able to produce a single uniform gradient from the four mean generalization gradients reported by Guttman and Kalish (1956). The transformation changed the relative interstimulus distances between wavelengths, resulting in large separations between 490 and 510 nm, 530 and 550 nm, and 600 and 620 nm. These spectral regions correspond to those where Guttman and Kalish (1956) found the least amount of stimulus generalization—that is, where hue discriminability was apparently very good. Wright's (1972) psychophysical data certainly support Shepard's earlier conclusion that his unique interstimulus spacing probably had some important psychological significance.

INHIBITORY STIMULUS CONTROL

When some aspect of a stimulus previously correlated with reinforcement is varied during a generalization test, a decremental gradient with maximum height at $S+$ is usually obtained. Such gradients are often called excitatory gradients (Hearst, 1968; Terrace, 1966b), as they are thought to reflect excitatory stimulus control by $S+$. The question arises quite naturally whether corresponding incremental gradients can be obtained around a stimulus correlated with extinction ($S-$). Such gradients might be the outcome of inhibition generated by responding without reinforcement, and would thus be called inhibitory gradients (Hearst, 1968; Terrace, 1966b), presumably reflecting inhibitory stimulus control. In this section, we review the methods used to study generalization around a stimulus correlated with extinction, summarize some important results, and discuss problems of interpretation.

The basic method of testing for inhibitory stimulus control is the same as that developed by Guttman and Kalish (1956). Values on the test dimension are presented repeatedly, usually (but not always) without reinforcement, and responding at each test value is measured. The training procedure preceding the generalization test, however, differs from that typically used in the study of excitatory stimulus control. Responding must initially be reinforced in the presence of some stimulus other than $S-$ in order to generate a level of responding against which control by $S-$ may be assessed. (The symmetrical operation, extinguishing responding to a value other than $S+$, is not required in assessments of excitatory stimulus control.) There are various ways to generate this baseline response level: by reinforcing responses to a range of values on the test dimension, by reinforcing responses to a stimulus orthogonally related to the test dimension (interdimensional discrimination training), and by reinforcing responses to another value on the same dimension (intradimensional discrimination training). We will describe each of these methods in turn. In our discussion, we will refer to stimulus control gradients around $S-$ as "incremental" or "negative" rather than "inhibitory," in order to "separate" empirical results from theoretical connotations.

Gradients of Extinction

Honig (1961) was the first to study generalization of extinction in pigeons. In his experiment, a baseline level of responding was first established by reinforcing key pecking at 13 different spectral values, ranging from 490 to 630 nm. When response rates were stable, two groups were extinguished in the presence of 570 nm, one for 20 min, the other for 40 min. Gradients following extinction were then obtained in the manner of Guttman and Kalish. These gradients were bowl-shaped with minima at 570 nm. Furthermore, the reduction in responding at that value was greater for the group receiving the longer period of extinction. A control group which had not received extinction provided flat gradients.

Honig's pigeons obviously learned the association between a particular spectral value and extinction. Their gradients were much shallower, however, than those obtained following simple excitatory training. The reason for this, Honig suggested, was that extensive reinforcement with a large number of stimuli prior to extinction had made the subjects less sensitive to differences in spectral value, thus flattening the gradients. Honig also suggested that the incremental gradients did not necessarily indicate inhibitory stimulus control by *S*−. Extinction may simply have reduced excitation to a greater degree at that value than at others. This distinction between reduced excitation and inhibition has been carefully developed by Hearst, Besley, and Farthing (1970) and Hearst (1972), and we will return to it shortly.

Blough (1975) reported extinction effects similar to Honig's in a steady-state procedure. A baseline was initially established by scheduling intermittent food reinforcement and continuous secondary reinforcement for responding to 25 closely spaced stimulus values. When responding had stabilized at each value, additional unreinforced trials were presented at the central value (597 nm). This procedural change had three effects. First, there was a general increase in responding to most of the baseline stimuli. Second, responding was reduced to *S*− and to the two adjacent values on either side of *S*−, producing a narrow negative gradient. Third, responding was elevated at values next to the gradient depression, creating shoulders. All of these effects can be seen in Figure 7. Several features of this gradient differ from Honig's, perhaps reflecting the fact that it was not the outcome of a pure process of extinction. Indeed, the shoulders effect and the overall increase in baseline responding are typical of the effects of discrimination training. They represent a form of the peak shift and behavioral contrast, phenomena discussed later in some detail. Furthermore, the narrow range of reduced responding suggests that all of the baseline stimuli were functionally positive.

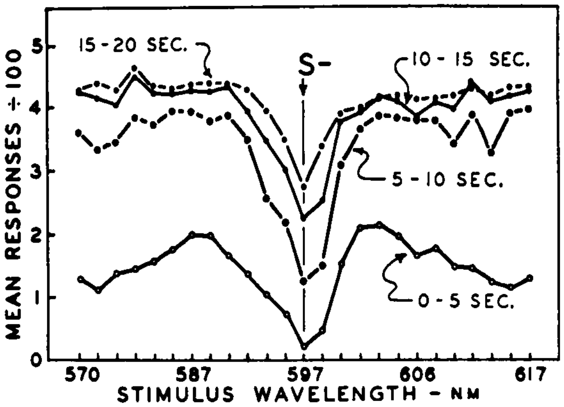

Fig. 7. Gradients of partial extinction obtained from a single pigeon across successive quarters of 20-sec trials. Secondary and primary reinforcement were scheduled at all test values with additional unreinforced trials at *S*−. Responding increased as each test trial elapsed, although gradient shoulders were evident only during the first two quarters of each trial. (From D. Blough, 1975).

Incremental Gradients: The Interdimensional Procedure

Interdimensional discrimination training—so named by Switalski, Lyons, and Thomas (1966)—provides a method for obtaining negative gradients without the necessity of reinforcing responding to stimuli which lie on the test dimension. In this procedure, subjects are initially trained to discriminate between an *S*+ and an *S*− selected from different stimulus dimensions. Generalization gradients around *S*− are then obtained by varying the *S*− characteristic which is not shared by *S*+. Since this characteristic is assumed to be orthogonal to that defining the *S*+ dimension, any influence of the latter upon responding should be equal across test values. The interdimensional procedure is also useful because the experimenter does not have to rely upon bouts of continuous extinction in order to obtain negative gradients (Honig, 1961), nor must he or she deal with the interpretation problems which arise if responding during *S*− had been previously reinforced (Blough, 1975).

An experiment by Honig, Boneau, Burstein, and Pennypacker (1963) illustrates this procedure. One group of pigeons (line-positive group) received discrimination training between a black vertical line on a white background (*S*+) and the white background alone (*S*−). For the other (line-negative) group, the white background was *S*+ and the vertical line was *S*−. After training to asymptote on their respective discriminations, the birds in both groups received a generalization test with different line orientations. The line-positive group produced an orderly, decremental generalization gradient with a peak at the vertical *S*+. The line-negative group, on the other hand, produced an incremental gradient with a minimum at the vertical *S*−. This negative gradient was shallower than the corresponding positive gradient, although when plotted in relative terms, it was rather similar in slope. Generalization effects like these were also obtained by Jenkins and Harrison (1962) and Schwartzbaum and Kellicutt (1962) using slightly different paradigms. In the latter study, for instance, rats produced incremental response gradients with different tone frequencies when the presence of a particular tone had served as *S*− and its absence as *S*+ during interdimensional training.

Beale and Winton (1970) studied stimulus control by *S*− in a procedure where pigeons could control the duration of each discriminative stimulus. Pecks at a separate control key alternated between a blue light (*S*+) and a black vertical line superimposed on the blue surround (*S*−). The pigeons quickly learned to spend most of the time during training in the presence of *S*+. Two separate generalization tests, each conducted with six different orientations of the black line, followed discrimination training. The control key was covered in one test, with each trial lasting 1 min and response rate providing the only measure of stimulus control. In the other test, the control key was available so the birds could proceed from any test stimulus to the next programmed value. There was no time limit on any trial, and both stimulus duration and response rate were dependent variables. The results from control key test are shown in Figure 8. The gradients of stimulus duration and of total responses are incremental, although the former is steeper and more orderly than the latter. The gradient based on response rate is also incremental and similar in form to the one obtained from the test without the control key (not shown in the figure).

A comparison between the slopes of negative and positive gradients following symmetrical training procedures has been of some interest in this line of research. Honig et al. (1963) argued that both types of gradients were similar in slope, but the balance of the literature suggests that negative gradients are, in fact, flatter than positive gradients (cf. Boneau & Honig, 1964; Honig & Beale, 1976; Jenkins & Harrison, 1962). Jenkins (1965) provided an interesting interpretation for this asymmetrical result, which we outline below.

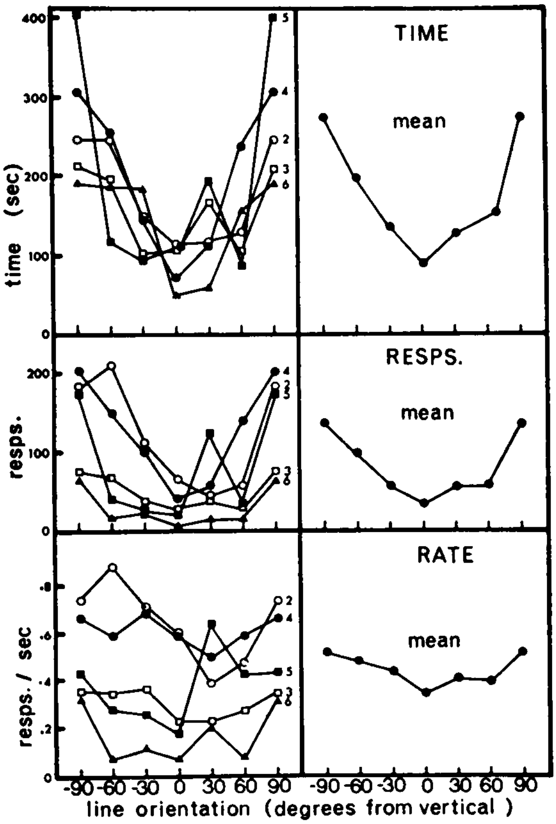

Fig. 8. Individual and mean negative gradients obtained from pigeons who could advance from one test stimulus to the next by means of a separate control key. Three measures of stimulus control are shown: time spent in each stimulus, mean number of responses, and response rate. The steepest gradient is that based on stimulus duration. (From Beale and Winton, 1970).

Designate the key peck response as R_1. All other types of behavior can then be designated as $\overline{R_1}$. This class of "other behavior" or "not key pecking" may, however, have some identifiable (although not recorded) subclass of behavior within it—for instance, pacing back and forth at the sides of the chamber. Assume that pacing is an unrecorded but identifiable subclass and designate it R_2. All other types of behavior besides key pecking and pacing can now be called $(\overline{R_1 + R_2})$. Consider now the generalization test following conditioning of R_1 to a specific $S+$. When some feature of this stimulus is varied, the typical result is an orderly decrease in R_1, which may reflect increases in R_2, or in $(\overline{R_1 + R_2})$, or in both. But no matter how responses are distributed into these latter two classes of behavior, all losses are summed to give a net decrease in R_1, the measured response.

Stimulus control by $S-$ may also be analyzed in this way. Assume that an identifiable subclass of "other behavior," like pacing (R_2) has been conditioned to $S-$. (Such behavior might be reinforced by the reappearance of $S+$.) During a generalization test, variations in some feature of $S-$ result in losses of behavior (R_2) controlled by $S-$. These losses can be reflected as increases in R_1 (key pecking), or as increases in $(\overline{R_1 + R_2})$, or both. However, only R_1 is measured and used to plot the generalization gradient. If behavior lost from R_2 were to transfer completely into R_1 (normally key pecking), then the negative gradient around $S-$ would be comparable to the positive gradient around $S+$. However, any losses from R_2 into $(\overline{R_1 + R_2})$ reduces the number of responses which constitute the negative gradient. Consequently, it will be flatter than the positive one, in spite of symmetrical test procedures.

The Post-Discrimination Gradient: Effects of Intradimensional Training

A test of Spence's theory. Perhaps the oldest method for establishing stimulus control is to provide discrimination training between two stimuli which lie on the same dimension. This procedure often produces transposition when an appropriate choice test is carried out (Honig, 1962). In other words, if the subject is offered two stimuli, one of which is the original $S+$, the other a novel stimulus that lies beyond $S+$ with respect to the negative training value, the latter is often chosen. The phenomenon of transposition was not easily encompassed by traditional theories of discrimination learning. Spence (1937), however, proposed an ingenious explanation based upon hypothetical gradients of generalization. He suggested that, in the course of discrimination training, $S+$ acquires excitatory properties and $S-$ acquires inhibitory properties, and that these properties generalize decrementally to other stimuli along the same dimension. Spence argued that the resulting theoretical gradients would summate algebraically to produce a postdiscrimination gradient (PDG) from which transposition could be predicted. His theory also led to a number of other predictions which were subject to experimental test once the proper procedures were developed for the purpose.

(1) Since the inhibitory gradient is subtracted from the excitatory one, the height of the PDG should be reduced relative to a gradient obtained after training with $S+$ alone.

(2) The theoretical PDG should result in a mode, or peak, which is not at $S+$, but at a value displaced from $S+$ in a direction away from $S-$. (Spence had used this prediction of peak shift to explain transposition.)

(3) The PDG should be steeper between $S+$ and $S-$ than a gradient obtained after training with $S+$ alone.

(4) The size of the peak shift should be inversely related to the $S+/S-$ difference during discrimination training.

Hanson (1959) evaluated these predictions by testing four groups of pigeons for spectral generalization after discrimination training between different pairs of hues. All groups were trained with the same positive stimulus, 550 nm, but differed from one another with respect to the stimulus associated with extinction. The values of $S-$ were 555, 560, 570, and 590 nm. The generalization gradients obtained from these four groups and the gradient for a control group trained with 550 nm alone are shown in Figure 9.

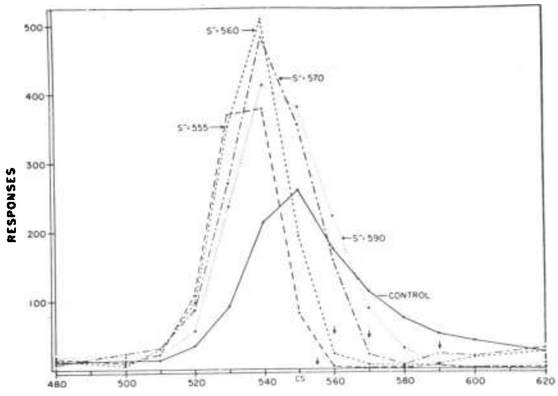

Fig. 9. Mean gradients obtained from four groups receiving discrimination training between 550 nm ($S+$) and one of four $S-$ values, as labeled on each gradient. Generalization data from a control group trained only with 550 nm are also shown. (From Hanson, 1959).

The most obvious characteristic of the PDGs is that they are, without exception, higher and steeper than the control gradient. Each PDG also shows a peak shift, the magnitude of which is roughly proportional to the $S+/S-$ difference during discrimination training. Three predictions based on Spence's gradient-interaction theory are therefore confirmed. These concern the steepness of the PDG, the shift in the gradient peaks, and the amount of peak displacement. The only unconfirmed prediction concerns the height of the PDG. Spence's theory predicts that the PDGs should be lower in height than the control gradient because inhibition should reduce responding along the entire gradient. There are two possible explanations for this discrepancy. One is that Hanson's control group, reinforced at 550 nm alone, was given less training at that value than were the experimental groups. A more probable explanation, however, is based on behavioral contrast, a phenomenon of discrimination learning not studied until the early 1960's. Behavioral contrast refers to an increase in response rate during $S+$ which accompanies a decrease in rate during $S-$ (Reynolds, 1961a). A contrast effect probably carried over from training to testing, producing the high peaks in the PDGs (see Friedman & Guttman, 1965, and Terrace, 1968, for further research and discussion of this matter).

The negative peak shift. A natural question arising from Hanson's experiment concerns the negative peak shift—that is, minimum responding at a value beyond $S-$ in a direction away from $S+$. Hanson (1959) did not observe a negative peak shift, but this may have been due to his stringent discrimination criterion: an essentially zero response rate at $S-$. Clearly, rates below zero cannot be observed. Guttman (1965), however, reported a study where response rates were elevated before testing for a negative peak shift. Guttman initially reinforced responding equally to 19 different spectral values, ranging from 510 to 600 nm. Discrimination training between 550 and 560 nm was then carried out to a criterion less stringent than Hanson's. All of the original training values were then presented in a generalization test in extinction. Guttman observed both a negative and positive peak shift, as seen in Figure 10.

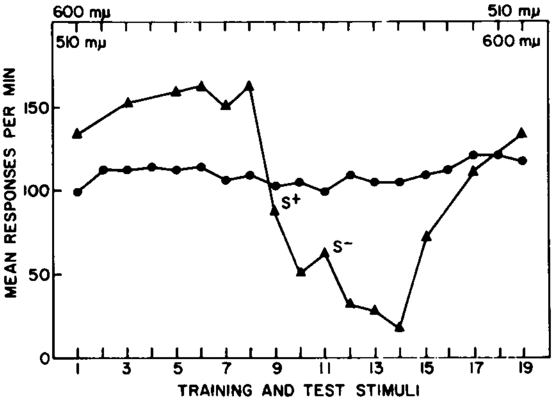

Fig. 10. Spectral generalization gradients averaged across six pigeons following two-part training consisting of equal reinforcement at all test values and subsequent discrimination training between 550 and 560 nm. The former value was $S+$ and the latter was $S-$ for four birds; for two others, the valences were reversed. Both positive and negative peak shifts are evident. (From Guttman, 1965).

Blough (1973, 1975) provided additional evidence for the negative peak shift in two more recent experiments. In one (Blough, 1975), he used the steady-state testing procedure previously described. Conditioned reinforcement was provided for responding to each of 25 spectral values, ranging from 570 to 617 nm. The conditioned reinforcer was a steady gray light for the 13 shortest wavelengths and a flickering light for the 12 longest wavelengths. These lights were associated with different probabilities of primary (food) reinforcement. In the final training condition, primary reinforcement was four times more frequent for responding to the longer wavelength values than to the shorter values. The steady-state data reflected these contingencies nicely: response rates were low with the shorter wavelengths, high with the longer wavelengths, with a smooth transition between the two. Furthermore, the decrease in responding at the negative values was accentuated at the transition between the high- and low-reinforcement conditions. This created a trough in the short-wavelength region of the gradient.

In a second experiment, Blough (1973) used a different method of testing to assess the negative characteristics of the stimuli around $S-$. Two groups of pigeons were run as a systematic replication of Hanson's study. One received positive training with 550 nm displayed on the key. The other received discrimination training between that value and 559 nm as *S*−. Generalization gradients for these groups were then obtained in extinction with test values spaced 9 nm apart. These gradients were very similar to those Hanson observed for his control group and for the group run with 560 nm as *S*−. Blough's third group was of particular interest. It was run on the same spectral discrimination, but had an alternate key illuminated with a white diamond available. Pecks at this key were reinforced when 559 nm appeared on the "spectral" key. Subjects trained in this way provided two separate gradients during their test. One was obtained from the spectral key, and showed a peak shift to about 540 nm. The other gradient was obtained from the alternate key. The birds responded substantially to this key in the presence of *S*−, but even more so when values further removed from *S*+ were displayed. The mean peak appears at about 575 nm. Thus, although values in the vicinity of *S*− failed to generate a negative peak shift on the spectral key, a corresponding positive peak shift was observed for those responses directed to the alternate key. This enhanced alternate responding indicates that values greater than 559 nm were more negative than *S*−.

Derivation of the postdiscrimination gradient from positive and negative gradients. One final question generated by the postdiscrimination data is whether the PDG with its peak shift can be derived from the summation of positive and negative gradients. This was the theoretical approach taken by Spence (1937). Hanson (1959) approached the issue by deriving hypothetical positive and negative gradients from his postdiscrimination data. His hypothetical gradients did not, however, conform to Spence's assumptions, so Hanson rejected a gradient-interaction analysis. Rilling (1977), on the other hand, reported a mathematical modelling analysis by Klein which showed that at least one of Hanson's empirical PDGs could be derived from hypothetical positive and negative gradients whose forms were reasonably consistent with those suggested by Spence. In his derivation, Klein assigned a greater absolute value to every point on the positive than on the negative gradient. This insured that all points on the derived gradient would be greater than zero. Klein also made the slope of the negative gradient steeper than the slope of the positive gradient in the vicinity of *S*+, a manipulation which guaranteed a positive peak shift. A small negative peak shift was also generated by these conditions. Interestingly, Spence's model does not predict a negative peak shift, although we have seen that such shifts are observed given a sufficiently high baseline.

Another approach to the derivation of PDGs is to combine positive and negative gradients obtained empirically. This has become possible with the interdimensional procedures described above. The PDGs derived from this method can then be compared with PDGs obtained following intradimensional training. Hearst (1968, 1969) used precisely this strategy in three separate experiments. One group of pigeons in each experiment provided a positive gradient following interdimensional training with a vertical line as *S*+ and a homogeneous white background as *S*−. A second group provided a negative gradient following training with the white background as *S*+ and a line as *S*−. The *S*− line differed from vertical in separate experiments either by 90°, 60° counterclockwise, or 30° clockwise. A third group received intradimensional training between vertical as *S*+ and the line orientation used in the same experiment to provide the negative gradient. Hearst then compared the PDGs from the intradimensional groups with the PDGs derived by adding together the independently obtained positive and negative gradients from each experiment. He found that the derived PDGs provided a good approximation to the empirical PDGs. His results are shown in Figure 11.

Hearst's method is elegant, but at least two aspects of his data analysis are questionable. First, in order to derive the PDGs from the empirical gradients, all response values along the negative gradients were adjusted to be less than or equal to zero. Yet, response rates were never less than zero. Second, in order to compare the derived with the empirical PDGs, all values along the former could be no less than zero, so a constant was added to each derived value. Agreement between gradient pairs was then measured by comparing corresponding indices of gradient asymmetry: the ratio of the total number of responses on the right side of *S*+ (clockwise tilts) to the total number on both sides. Unfortunately, the linear transformation of the derived data was not a ratio-preserving one, so it is difficult to interpret the purported agreement in gradient forms.

Finally, we should point out that Hearst predicted a peak shift in one of his experiments, where *S*— was 30° tilted clockwise, but did not obtain it. This is, in fact, a critical point of comparison between the predicted and the obtained data. After all, Spence's theory was invoked specifically to deal with the peak shift. Marsh (1972), however, noted that the peak shift is hard to obtain on the dimension of line orientation (but see Bloomfield, 1967) and, consequently, replicated Hearst's study using spectral stimuli. In his study, 550 nm was *S*+, 560 nm was *S*—, and the extradimensional stimulus (corresponding to the absence of the line in Hearst's study) was a white key. His treatment of the data was essentially similar to Hearst's method and can be criticized on the same grounds. However, Marsh did obtain a peak shift to 540 nm.

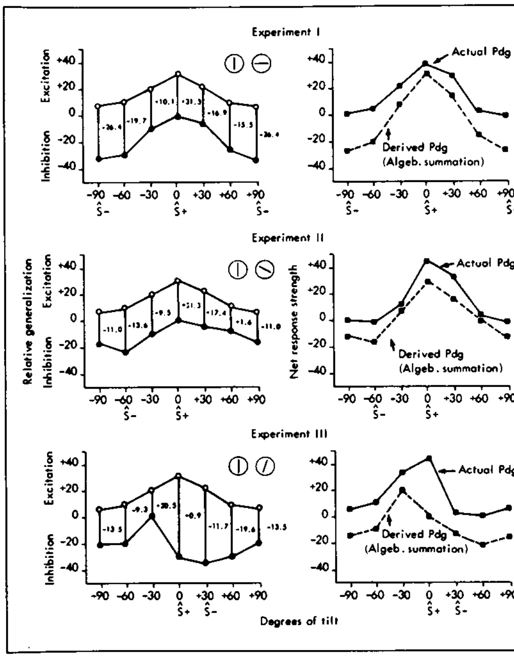

Fig. 11. Derivation of post-discrimination gradients from independently obtained positive and negative gradients. The latter were converted to relative values and are shown in the left panels. The differences between relative gradients are plotted as derived PDGs in the right panels, where they can be compared to empirical PDGs. (From Hearst, 1969).

Negative Stimulus Control and the Concept of Inhibition

The definition of inhibitory control. The rubric of inhibitory stimulus control is a convenient and generally accepted term for the description of the effects associated with *S*— in discrimination training, especially incremental gradients and the peak shift. However, the use of inhibition as an analytic concept involves several problems, one of which is whether a negative gradient constitutes evidence for an inhibitory process. For example, the gradients of extinction obtained by Honig (1961) are similar in form to negative gradients following interdimensional discrimination training. Yet, Honig's gradients were obtained after massed extinction, which produces neither behavioral contrast nor the peak shift (Honig et al., 1959). Furthermore, Terrace (1966b) obtained incremental gradients following interdimensional training only if subjects made frequent responses to *S*— (errors) while acquiring the discrimination. If the discrimination was learned without errors (see Terrace, 1966a), the interdimensional gradient was both low and flat. Terrace argued that incremental gradients from "errorful" subjects were evidence for inhibition, whereas the flat gradients from "errorless" subjects indicated no inhibition. Deutsch (1967) later disagreed with this interpretation claiming that low flat gradients indicated substantial inhibition, since virtually no responses occurred during testing.

In view of this kind of controversy, Hearst, Besley, and Farthing (1970) suggested that the operational criteria for inhibitory (and for excitatory) stimuli be independent of the dimensional control they exert over responding. An inhibitory stimulus would be defined as one which reduces the level of responding maintained in some other independent fashion. In the ideal experiment, *S*— from an incremental or a postdiscrimination gradient would be combined with a positive extradimensional stimulus with which training had been carried out independently. The putative inhibitory stimulus should reduce responding to this positive value. This result would naturally require a control condition showing that the reduced responding to the novel compound was not simply the result of generalization decrement. The same *S*— would thus be presented as part of a similar compound in a group where there was no prior association with extinction.

Clearly, the demonstration that *S*— is inhibitory is not a trivial matter, although it tends to be neglected in practice. Negative gradients may simply be the result of reduced excitation, so the conclusions drawn from this type of test result should be taken with caution. On the other hand, it appears generally accepted that the peak shift and behavioral contrast are reasonably good indices of inhibitory stimulus control. It may be wise nonetheless to supplement these latter results with an independent assessment of the properties of *S*−.

With these cautions in mind, we may proceed to the more difficult question of the training conditions that are necessary and sufficient to generate inhibitory control. Normally, discrimination training involves several confounded aspects and effects: (1) *S*+ and *S*− are juxtaposed in some temporal sequence, normally irregular alternation. (2) The two stimuli are associated with different reinforcement schedules, one of which is often extinction. (3) Responding to *S*− is usually reduced from some initial level. (4) Response rates to *S*+ and *S*− differ by the end of training. Furthermore, since procedural variables are typically confounded with their effects, a clear analysis of the conditions underlying inhibitory stimulus control is far from simple.

Juxtaposition of S+ and S−. Honig, Thomas, and Guttman (1959) initially trained pigeons to peck at 550 nm and then gave them a session of massed extinction at 570 nm without further reinforcement at 550 nm. A subsequent generalization test failed to provide evidence for a peak shift, although the overall level of responding was clearly reduced relative to a control group that did not receive extinction prior to testing. Similarly, Weisman and Palmer (1969) failed to obtain a negative gradient following massed extinction in an interdimensional paradigm. These studies demonstrate that inhibitory control is not simply generated by the correlation of a dimensional stimulus value with extinction. Other features of the discrimination procedure are apparently critical.

Rosen and Terrace (1975) replicated the findings of Honig et al. (1959) and then proceeded to a set of analytical studies in which one of three treatments followed massed extinction: (1) *S*+ was presented for 3 min with responding reinforced according to the same schedule used in training; (2) *S*+ was presented for 3 min without reinforcement; (3) the response key remained black for 3 min, but food was periodically delivered independently of any responding. Each of these three treatments generated a peak shift for intradimensionally trained subjects and a negative gradient for interdimensionally trained subjects, although the latter effect was not as consistent nor as orderly as the former. Their findings indicate that a certain temporal sequence of reinforcement and extinction is necessary to generate inhibitory stimulus control. Specifically, it appears that a period of reinforcement or one previously associated with reinforcement must follow extinction for inhibition to occur at *S*−. The data also indicate that any one element comprising the reinforcement period is sufficient to produce the effect and that the temporal alternation need only occur once.

Contrasting conditions of reinforcement. Extinction is generally the normal negative condition in discrimination training. Traditional S-R theory assumed that reinforcement leads to an excitatory process, whereas extinction leads to an inhibitory one. However, shortly after Hanson carried out his original study, Guttman (1959) reported a peak-shifted PDG following discrimination training between two stimuli signaling different schedules of positive reinforcement. Pigeons were trained on a hue discrimination task similar to Hanson's but with different positive schedules in effect from the outset of training: VI 1-min in the presence of 550 nm, and VI 5-min in the presence of 570 nm. After 12 training sessions, Guttman tested each subject with the same generalization values used by Hanson (1959). The mean gradient obtained from this test was remarkably similar to Hanson's corresponding PDG. In fact, when the gradients were normalized to provide equal values at *S*+, they were almost identical. Guttman's finding indicates that extinction is not required to make a stimulus functionally negative; it need only be associated with a less favorable reinforcement schedule.

Blough's (1975) steady-state study supports Guttman's findings. In Blough's experiment, different probabilities of primary reward were associated with two adjacent ranges of spectral values. The resulting steady-state gradient showed both a positive shoulder and a negative trough. Again, these effects were obtained even though all stimuli were associated with positive reinforcement.

Reduction in responding to S— and behavioral contrast. Free-operant discrimination learning normally involves a substantial reduction in responding to S—, frequently accompanied by an increase in responding to S+ (behavioral contrast). Terrace has suggested that these features of discrimination training are important factors in inhibitory stimulus control. One set of studies (Terrace, 1968; Weisman, 1969) is based upon the Guttman (1959) experiment. Terrace (1968) trained two groups of pigeons on an intradimensional hue discrimination task much like Guttman's. Prior to the discrimination phase, however, both groups received equal reinforcement for responding to both hues. One group was initially trained with equal VI 1-min reinforcement schedules, the other with equal VI 5-min schedules. The final discrimination schedule for both groups was multiple VI 1-min VI 5-min. Thus, the first group suffered a downshift from VI 1-min to VI 5-min on S—, whereas the other enjoyed an upshift from VI 5-min to VI 1-min on S+. The groups receiving the downshift showed a reduction in rate to S—, considerable behavioral contrast, and a peak shift during postdiscrimination generalization tests. Most birds in the upshift group showed neither effect.

Weisman (1969) reported similar differential effects on negative gradients. In his experiment, equal VI 1-min or VI 5-min schedules were initially programmed during each of two discriminative stimuli: a green hue and a vertical line superimposed on the green surround. Pigeons were then shifted to a discrimination procedure where VI 1-min reinforcement was scheduled during green and VI 5-min reinforcement was scheduled during green plus line. Negative line-tilt gradients were observed only in the group initially receiving equal VI 1-min training.

These experiments are interesting, but they are far from conclusive. For one thing, Guttman (1959) obtained the peak shift in spite of the fact that there was never any reduction in response rate to the stimulus associated with the poorer reinforcement schedule. Rates of responding to both stimuli increased during training, albeit at different rates. Thus, his experiment shows that a reduction in responding is not necessary in order to obtain inhibitory control by S—. Nor is it sufficient. If it were, then massed extinction, which certainly reduces response rates to S—, should produce negative gradients and the peak shift, which it does not. Behavioral contrast may thus suggest the presence of inhibitory stimulus control, but apparently it is not a determining factor.

Terrace's (1963a,b) studies on errorless discrimination learning also address this question, since there was no reduction in responding to S— (no responses to that value occurred during training). The principal results from the errorless tasks are an absence of peak shift (Terrace, 1964) and a failure to find incremental gradients (Terrace, 1966b). We have noted that a failure to find incremental gradients does not necessarily imply the absence of inhibitory stimulus control—responding may be too suppressed to show any dimensional effects. The failure to find a peak shift is more convincing (but see Terrace, 1966c). Nonetheless, Rilling (1977) has pointed out that the absence of errors is confounded with those training conditions that lead to the errorless discrimination phenomenon. Consequently, we do not know whether the introduction of S— early in training (and at a reduced level of duration and/or intensity), or the lack of unreinforced responses to S—, eliminates potential inhibitory control.

It is not easy to conclude that a single set of conditions is necessary or sufficient to generate inhibitory stimulus control. Terrace has suggested that such control is exercised by stimuli that are functionally "negative," and errorless training serves to make an S— "neutral" rather than negative. This description probably adds little to the empirical findings on which it is based. In any case, no independent set of criteria seems to be available which will identify a "neutral" stimulus independently of its failure to generate peak shift and related phenomena. It is also difficult to see why a "neutral" stimulus resulting from massed extinction should suddenly turn negative when this procedure is followed by a few response-independent reinforcers delivered in the absence of any designated discriminative stimulus.

The Resistance-to-reinforcement Test

We have seen that the usual generalization test for inhibitory stimulus control may be insensitive because of floor effects: low levels of responding which preclude detection of any dimensional effects. An alternative measure of inhibitory control would therefore seem desirable. Stimulus duration is one such alternate measure, because it may potentially enhance small dimensional effects in gradients based upon response rate (cf. Beale & Winton, 1970). Another alternative that has yet to be discussed and that has received some recent attention is the resistance-to-reinforcement test proposed by Hearst et al. (1970). Its use has been reported by Zentall, Collins, and Hearst (1971), Karpicke and Hearst (1975), and Rilling, Caplan, Howard, and Brown (1975).

The resistance-to-reinforcement test is identical in every respect to the procedure for obtaining negative gradients, except that responding is reinforced at each test value. The test is modelled after the retardation-of-learning test in classical conditioning (Rescorla, 1969) and its rationale is as follows. If *S*− has acquired inhibitory properties via its association with extinction, then subsequent correlation with reinforcement should produce slower response acquisition than if the stimulus had either been novel, uncorrelated with reinforcement, or previously correlated with reinforcement. Inhibitory control would result in a response gradient whose height (at least at the *S*− value) would be lower than gradients obtained after training with any of the control conditions.

An experiment by Zentall, Collins, and Hearst (1971) illustrates this and the usual extinction procedure. Zentall et al. trained two groups of pigeons on an interdimensional discrimination between the presence (*S*−) and the absence (*S*+) of a vertical line. One group had been previously trained on the reverse discrimination (line-present was *S*+, line-absent was *S*−). Generalization tests with different line orientations were then given in extinction and, afterwards, with reinforcement at all dimensional values. Incremental response gradients with minima at *S*− were obtained from both groups in both tests, although the overall rate of responding was higher during the reinforcement test than during the extinction test. Furthermore, group differences emerged as a function of prior reinforcement history. The group which had reversal training responded more frequently during testing than the group without such prior training. In addition, repeated resistance-to-reinforcement tests revealed decremental response gradients around the vertical line in the reversal group. The nonreversal group continued to provide incremental gradients.

Rilling, Caplan, Howard, and Brown (1975) demonstrated the importance of elevated response levels for detection of inhibitory stimulus control. Three groups of pigeons were trained through autoshaping to distinguish between a green hue as *S*+ and a line on a white background as *S*−. This particular discrimination insured very few responses to *S*−. All groups were then tested with different line orientations. One group had a compound test in extinction: each orientation was superimposed on the green hue. The other groups received resistance-to-reinforcement tests. The lines were compounded with *S*+ for one group, but not for the other. The resistance-to-reinforcement groups provided shallow negative gradients which were parallel to one another and separated by an amount presumably attributable to the presence vs. absence of the green surround. Furthermore, the negative gradients for some birds appeared only after reinforcement had raised the overall level of responding. The group tested in extinction provided a low flat gradient, in spite of the presence of the green background. Clearly, this floor effect prevented the appearance of inhibitory effects along the test dimension.

There is one major disadvantage to the resistance-to-reinforcement test which should be pointed out. If little or no responding occurs to *S*− (or stimuli adjacent to it) during the test, while relatively higher rates of responding occur to test values further removed, the negative gradients may simply be due to unequal reinforcement density across the test continuum. Specifically, reinforcement may be less frequent at or near *S*−, thus contributing to a gradient depression in this region. Hearst et al. (1970) and Rilling et al. (1975) claim that this has not been a problem in their experiments. Nonetheless, the potential for confounding is well illustrated by Karpicke and Hearst (1975), who reported that explicit response shaping was necessary for some birds during the initial phases of testing. With this caution in mind, however, the resistance-to-reinforcement test promises to be a standard assay for inhibitory stimulus control in future research.

ATTENTIONAL PROCESSES AFFECTING THE GENERALIZATION GRADIENT

Postdiscrimination generalization gradients obtained after intradimensional training are quite steep, not only between *S*+ and *S*− but also on the other side of the shifted peak. This apparent increase in slope is usually confounded with the increased height of the PDG resulting from behavioral contrast (Hanson, 1959). Nonetheless, the effects of discrimination training appear to be more general than the local changes associated with facilitation of responding around *S*+ and suppression around *S*−. The study of such general changes has been profitably carried out with the use of interdimensional training procedures. We have just seen that the interdimensional design permits an assessment of inhibitory stimulus control uncontaminated by excitatory effects. The same approach helps to separate the general effects of discrimination training from the specific. These general effects are commonly identified as "attentional" and are the topic of this section.

Jenkins and Harrison (1960) were perhaps the first to describe the general effects of discrimination training upon the slope of the generalization gradient. In their experiment, two groups of pigeons were trained to peck at a white key in the presence of a 1000-Hz tone. One group received simple acquisition (i.e., nondifferential) training to the tone; the other received discrimination training between the tone (*S*+) and its absence (*S*−). Following training, generalization tests were given along the dimension of tonal frequency. These tests showed that nondifferential training produced virtually flat gradients, whereas postdiscrimination gradients were decremental, very orderly, and peaked at *S*+. Similar effects due to training condition were found in a subsequent experiment in which two tonal frequencies, 450 and 2500 Hz, served as positive stimuli. A common interpretation of these results is that discrimination training between the presence and absence of the tone enhanced the subjects' attention to the tone, as manifested by decremental gradients along the relevant dimension. In the nondifferential condition, on the other hand, subjects attended to other, incidental cues present in the experimental environment.

A great deal of subsequent research has been devoted to effects like those reported by Jenkins and Harrison (1960) and generalization tests have frequently provided the critical assessments of stimulus control. Following the suggestion of Honig (1970), attentional effects can be identified as modulations of stimulus control which take the form of general differences in the slopes of generalization gradients. The differences in slope are due to specific experimental treatments carried out either in training or in testing and normally involve stimuli orthogonal to the test dimension. Honig's view encompasses the notion that, in one sense, attention denotes the very fact of stimulus control (Skinner, 1953): orderly changes in behavior result from changes in an effective stimulus dimension. It also suggests, however, that certain underlying processes make a particular dimension more or less effective in controlling behavior.

Two principal theories have been proposed to explain how these attentional processes operate. The first ascribes attentional effects to a selection among concurrently available stimuli. Attentional enhancement occurs, as it were, by default, in that competition from other stimuli is reduced in favor of the dimension on which the enhancement is observed. This particular view has been espoused by Sutherland and Mackintosh (1971, especially Chapter 7) and by Mackintosh (1977). It is also the account offered by Jenkins and Harrison (1960) for their results. The second is that attentiveness to all predictive stimuli is enhanced by discrimination training. According to this notion, attention generally increases with discrimination training and diminishes with equal reinforcement. Thomas and his associates (Thomas, 1970; Thomas, Freeman, Svinicki, Burr, & Lyons, 1970) have argued for this view. We will refer to both interpretations in our review of attentional phenomena in stimulus generalization.

Competition among Stimuli

We will begin our discussion with phenomena that reflect an inverse relationship between two independent stimulus dimensions with respect to their control over responding. Let us distinguish first between the stimuli from a *criterion dimension* and the stimuli from a *competing dimension.* The former are those that establish control in the course of training. The latter systematically alter control exercised by the former. Competing stimuli can be introduced at various stages in the experiment, so their effects are frequently labeled according to their point of introduction. For example, *overshadowing* may occur when the competing stimuli are presented simultaneously with the criterion stimuli during training. A related effect, *blocking*, may be observed when the competing stimuli are introduced prior to training and are maintained during training. Finally, *masking* becomes possible when competing stimuli are introduced only during the generalization test. Each of these phenomena can modulate attention to the criterion stimuli in its own way, so it is important to distinguish between their respective effects. Unfortunately, the competing stimuli in some experiments are present during both acquisition and testing, so it is not clear to what extent reduced control by the criterion dimension is due to overshadowing (or blocking), to masking, or to both.

Masking. Few experiments have been designed to study masking, although some unpublished results obtained by Honig and Gerry (Honig, Note 1) illustrate this phenomenon. Pigeons were first trained to discriminate between a high tone (2500 Hz) and a low tone (1000 Hz), during which time they pecked at a blue light. The high tone was *S*+ for two birds and *S*— for two others. The birds were then taught to discriminate between the presence (*S*+) vs. absence (*S*—) of three vertical lines on a white background. No tone was presented during the visual discrimination, although refresher sessions on the tone discrimination were periodically provided. During a subsequent generalization test, four different line orientations were presented under three separate conditions: in the absence of any tone (control condition), in the presence of the *S*+ tone (positive condition), and in the presence of the *S*— tone (negative condition). The corresponding gradients are shown in Figure 12. The control gradient was sharply sloped, with little responding to *S*—. The negative gradient, on the other hand, was relatively flat and markedly lower in height. The positive gradient was also rather flat, although elevated in comparison to the negative gradient. There was also considerable responding to the white key in the positive test condition, even though this stimulus was *S*— in the visual discrimination.

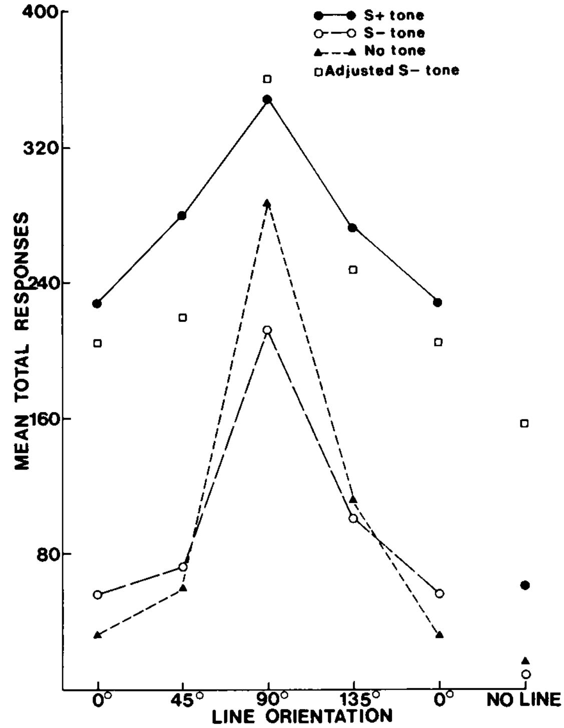

Fig. 12. Line-orientation gradients following training with a vertical line as *S*+ and the absence of a line as *S*—. Gradients were obtained in the presence of a positive or a negative tone from an independent discrimination, or in the absence of any tone. See text for explanation of unconnected symbols. (From W. K. Honig & J. Gerry, unpublished data).

One interpretation of these results is that the pigeons divided their attention between the lines and the tones during the positive and negative conditions. When they attended to the tones, the lines exercised no stimulus control. Conversely, when they attended to the lines, the tones exercised no control. This view predicts that the line orientation gradients ought to be roughly parallel and that any difference between them should only reflect the contribution of responses controlled by the tones. This contribution can be estimated as the mean difference between the line orientation gradients and the *S*— values for the two tone conditions. When this difference is added to each point along the negative-tone gradient, a good approximation to the positive-tone gradient is obtained, as shown in Figure 12 by the unconnected squares.

The Honig and Gerry study suggests that masking results from a division of processing time between stimuli from two modalities.

It also points out a danger in comparing gradients on the basis of normalized percentage values. Clearly, if the gradients obtained in the presence of the tones were normalized, the one from the negative condition would be much steeper. The processing-time analysis implies, however, that control by line orientation should be the same under both positive and negative conditions. One thing is certain: the gradient representing the greatest control by line orientation is that obtained in the absence of either tone.

Overshadowing. Unlike masking, overshadowing has been studied extensively. The first overshadowing study employing generalization techniques was conducted by Newman and Baron (1965). In their experiment, one group of pigeons was trained to discriminate a vertical line on a green background (*S*+) from the green background alone (*S*−). A second group was trained in a similar fashion, but *S*− was a red rather than a green field. Consequently, the line and color were redundant cues distinguishing *S*+ from *S*− in the latter group. Following training, a generalization test involving several different line orientations was given to both groups. The birds trained to discriminate solely on the basis of the vertical line provided orderly, decremental response gradients. The line-orientation gradients obtained from the subjects trained with the redundant color cue, on the other hand, were virtually flat, suggesting that color had overshadowed the line during training. The acquisition data support this interpretation. Pigeons trained with different backgrounds learned the discrimination faster and maintained higher performance levels than those for which the line was the only cue. We should note, however, that the colored backgrounds may have masked as well as overshadowed control by line orientation, since these cues were also present during testing.

Farthing (1972) separated overshadowing from masking effects in an experiment using similar stimuli (lines and colors) and a similar design. Pigeons were trained either to discriminate solely on the basis of a vertical line or on the basis of both line and color cues. Two types of generalization tests were then conducted. In one, only the line in one of several test orientations was presented on each trial. In the other, both the test line and the positive background color were presented together. Farthing's results replicated those of Newman and Baron (1965). The gradient obtained following discrimination training with the line as the only relevant cue was steeper than the gradient obtained when the background color had been redundant. Furthermore, these differences were evident even when the color was omitted during testing, although the corresponding gradients were greatly reduced in height and area. Thus, overshadowing was shown to emerge during training and to be independent of the presence of the overshadowing color during testing.

Taken together, the Newman and Baron (1965) and the Farthing (1972) studies suggest that the direction and degree of overshadowing is determined by the relative discriminabilities of the redundant cues. Color overshadowed the line in these experiments, presumably because it was the easier cue (cf. Lovejoy & Russell, 1967, and Chase & Heinemann, 1972). Miles and Jenkins (1973) extended these findings by demonstrating how redundant cues can overshadow one another in discrete-trial paradigms. For some birds (controls) in their experiment, the presence of a 1000-Hz tone was the only cue for discriminative responding. For others (experimentals), the light intensity projected onto the response key was redundant with the tone cue. The intensity on positive trials (tone present) was the same for all groups, whereas the intensity on negative trials (tone absent) varied between them. This manipulation allowed Miles and Jenkins to examine overshadowing as a function of the ease of the visual discrimination. Tests were then carried out for control by both the light and tone dimensions, with various light intensities combined factorially with the presence and absence of the reinforced tone.

The generalization data are presented in Figure 13. The higher gradient in each panel represents responding to the visual stimuli in the presence of the tone; the lower shows responses to the same stimuli in its absence. The slope of each gradient indicates stimulus control by light intensity, whereas the separation between gradients indicates control by the tone. For the control birds (upper left panel), the tone exerted the greatest degree of stimulus control: there was little responding in its absence, and the light intensity gradients were quite flat. For the experimental birds, tone control diminished (the gradients were less sep-

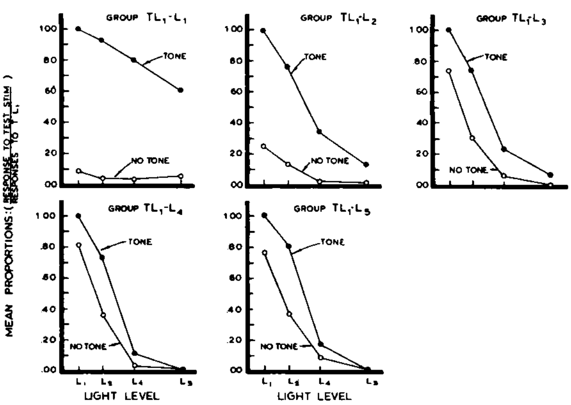

Fig. 13. Relative postdiscrimination gradients for light intensity obtained in the presence and the absence of a tone. *S*+ was always the compound of the tone (*T*) and the most intense light (L_1). *S*− was the absence of the tone and a light intensity which differed for different groups (L_1-L_5). (From Miles & Jenkins, 1973).

arated) and light control increased (the gradients became steeper) the greater the difference between the light intensities used in training. This demonstrates a reciprocal overshadowing relationship: light overshadowed tone when the light discrimination was easy, and vice versa when it was difficult. Furthermore, the slope differences between light intensity gradients appeared even when the tone was absent during testing, thus eliminating the possibility that the tone simply masked the light. It is possible, of course, that the light masked control by the tone, since no trials were run without the keylight. That test condition, however, would have eliminated responding entirely unless special training procedures had been used.

Studies such as these indicate that stimuli on a criterion dimension will be overshadowed by competing stimuli that also predict reinforcement. Wagner, Logan, Haberlandt, and Price (1968) have further demonstrated that the relative validities of the predictive cues will determine the degree to which each set gains control over behavior. Wagner et al. presented a light (criterion stimulus) on every conditioning trial with reinforcement probability set at 50%. One of two tones was also presented together with the reinforced light, but its relationship to reinforcement differed between groups. For one group in each study, the tones were differentially correlated with the presence and absence of reinforcement. For the other group, the tones were uncorrelated with reinforcement. Tests were then carried out with the light alone, with each of the tones alone, and with the light-tone combinations. For our purpose, the important result from each study was that the light gained much more control over responding when the tones were uncorrelated with reinforcement than when they were correlated.

Blough (1969b) and Heinemann and Chase (1970) have also examined stimulus relevance effects such as those reported by Wagner et al., but within a framework of generalization tests. In Blough's (1969b) experiment, pigeons were initially trained with all possible combinations of seven spectral lights and seven pure tones in a discrete-trial, steady-state discrimination procedure. Responding was reinforced during presentations of only one of the 49 possible combinations, extinction being in effect during the remaining 48. Thus, both spectral and auditory cues were relevant during discrimination training. After the spectral and auditory response gradients had stabilized, Blough reduced the relevance of each stimulus dimension in turn by holding either the light or the tone constant at its reinforced value. Differential training along the remaining dimension continued during this test phase. Blough found that the auditory and the spectral gradients sharpened in turn when the stimulus from the complementary dimension was held constant. Furthermore, gradients along the constant-stimulus dimension were considerably flattened upon return to bidimensional discrimination training, demonstrating a loss of stimulus control when these cues were irrelevant to discrimination performance. Although overshadowing may have produced these effects, one particular procedural detail precludes unqualified acceptance of this interpretation. When each *S*+ was made irrelevant, its associated reinforcement density was reduced. This could account for the flattening of the irrelevant gradients upon return to differential training (cf. Hearst, Koresko, & Poppen, 1964). The change in the effective reinforcement schedule would not explain, however, the sharpening of the "relevant" gradients with constant training along the "irrelevant" dimension.

Heinemann and Chase (1970) found results similar to those reported by Blough (1969b) in a between-groups design which controlled reinforcement density. The pigeons in their experiment were trained to make left vs. right choice responses on the basis of particular combinations of light and white noise intensities presented in a discrete-trial procedure. There were two intensity values for each dimension, yielding a total of four possible stimulus combinations. For some birds, the correct choice response on any given trial depended upon the stimulus intensities from both dimensions. For others, only light intensity (luminance) or white noise intensity signaled the appropriate choice response—the two values from the remaining dimension were uncorrelated with correct choices. Stimulus control by each dimension was then assessed in generalization tests during which eight luminance levels were paired factorially with eight noise intensities in extinction.

Heinemann and Chase found that the relative degree of control exerted by the light and white noise depended upon their predictive relationship to reinforcement during training. Both dimensions controlled choices when both had been relevant: the gradients of luminance and white noise intensity were clearly decremental. When only a single dimension had been relevant, that dimension alone controlled choice. Gradients along the irrelevant dimension were virtually flat, indicating that those stimuli had been overshadowed by the stimuli from the relevant dimension. Although neither stimulus dimension was tested in isolation, a simple masking interpretation cannot account for their findings, since all groups were tested under identical procedural conditions.

Blocking. A third mechanism by which stimuli can lose the control they would otherwise acquire is through a process called "blocking" (Kamin, 1968). Blocking refers to the finding that the amount of control acquired by one element of a compound stimulus is reduced or even eliminated by prior training with the other element in isolation. This phenomenon has been primarily studied within classical conditioning paradigms. Its effects on dimensional stimulus control, however, have also been assessed in a number of experiments using the generalization technique.

Johnson (1970), for example, trained pigeons on a free-operant discrimination between two stimulus compounds consisting of line orientations (vertical or horizontal) superimposed on monochromatic backgrounds (501 or 551 nm). One line-color compound served as $S+$, the other served as $S-$. A control group learned only the compound discrimination; a second blocking group learned to discriminate between the two line orientations prior to training on the compound task. Both groups then received generalization tests in extinction, during which five orientations of the line were combined factorially with five background wavelengths. The pigeons in the blocking group generally produced flatter wavelength gradients and steeper line-orientation gradients than the pigeons in the control group. Thus, pretraining along the line-orientation dimension not only decreased (blocked) wavelength control, but also increased orientation control.

Similar findings were later reported by Chase and Heinemann (1972). In their experiment, pigeons learned a two-choice discrimination between stimulus compounds of light and white noise intensity. Some birds were trained only on the compound task, whereas others received preliminary training with either the two light or the two white noise intensities, elements of the compound. Generalization tests following compound training generally revealed steeper gradients along the pretrained dimension and flatter gradients along the added dimension in the groups receiving prior training with one of the elements.

The results of these two studies and others like them (Mackintosh & Honig, 1970; vom Saal & Jenkins, 1970) can be most easily explained using the notion of stimulus relevance or stimulus validity mentioned in the previous section. In short, stimulus control by one set of cues can be blocked by another set if the former are no more (or no less) predictive of reinforcement than the latter to which it is added.

Interdimensional Discrimination Training and Generalization

We are now in a position to interpret the Jenkins and Harrison (1960) study in terms of the processes described above. For the groups that were trained nondifferentially, irrelevant stimuli (such as the light on the response key) may have overshadowed (or perhaps even blocked) control by the tone(s), since they predicted reinforcement just as well. They may have also masked control by the tone during testing. For the groups receiving discrimination training, however, the ambient stimuli were less predictive of reinforcement and so lost control over responding and the capacity to overshadow the tone (cf. Wagner et al. 1968). Consequently, generalization gradients based on tonal frequency became steeper.

Rudolph and Van Houten (1977) put this interpretation to a direct test. One group of birds in their experiment was trained like the Jenkins and Harrison nondifferential group. A second group also received nondifferential training but had the keylight gradually faded out over the course of training. Eventually, subjects in this group ended up pecking in the dark during tone presentations. Subsequent generalization tests revealed very shallow auditory gradients if the keylight had been present during training and steep gradients if it had been absent. This finding indicates that it was not any particular characteristic of the auditory stimulus in the Jenkins and Harrison (1960) study that caused its failure to acquire stimulus control. Rather, it had been overshadowed by the simultaneously present (redundant) keylight.

In the research described to this point, cues from the criterion and competing dimensions were presented together in training. Thus, one can envision a selection among simultaneously presented stimuli as a mechanism for overshadowing. Simultaneous training along two dimensions is not necessary, however, to obtain effects which reflect the relative predictive values of various cues. Switalski, Lyons, and Thomas (1966), for instance, presented a green keylight and a vertical white line on a dark background to pigeons on separate trials within a session. Under some conditions, the green key was $S+$ and the white line was $S-$. Under other conditions, equal reinforcement was provided in the presence of both stimuli. Spectral generalization gradients obtained in the absence of the white line were consistently steeper following discrimination training than following equal reinforcement. Lyons and Thomas (1967) modulated the slopes of gradients from individual subjects in a similar way by alternating discrimination training and equal reinforcement in advance of a series of tests.

These results are somewhat difficult to interpret in terms of overshadowing, since both color and line cues were presented successively rather than simultaneously. The notion can be preserved, however, if we recognize the role of (unspecified) background cues common to presentations of both. When the color and the line are equally correlated with reinforcement, neither is a better predictor of reinforcement than background cues. Consequently, the latter gain control over responding at the expense of color (and, presumably, line orientation). When color and line are differentially correlated with reinforcement, then background cues acquire less control, so the slope of the spectral gradient increases accordingly. Even this analysis must be accepted with some degree of caution, however, since the line in these experiments differed from the green key in color and in form. The effective discrimination may well have been intradimensional.

Extradimensional Discrimination Training and Generalization

The most general effects of discrimination training upon stimulus generalization can be studied if training is conducted along a dimension orthogonal to that on which stimulus control is acquired and tested. This is the *extradimensional paradigm.* During the extradimensional training phase, subjects learn to discriminate between two stimuli lying along a particular dimension. Responding is then shaped, either separately or concurrently, to a third stimulus (dimensional acquisition) from an orthogonal criterion dimension. Finally, a generalization test is carried out on the criterion dimension in the absence of the extradimensional stimuli.

Honig (1969) was one of the first to study attentional processes using this procedure. One group of Honig's pigeons was trained on a *true discrimination* between blue and green hues. Another group was trained on a *pseudodiscrimination,* in which reinforcement was available half of the time in the presence of both stimuli. Both groups were then trained to peck at three dark vertical lines on a white background. A generalization test with different line orientations followed. The finding of interest was that the gradient obtained from the true-discrimination group was considerably steeper than the gradient from the pseudodiscrimination group.

Variants of this procedure have been used to study the effects of extradimensional discrimination training when it precedes dimensional acquisition (Honig, 1969), is concurrent with such acquisition (Wagner, 1969), or is conducted afterwards (Honig, 1974; Turner & Mackintosh, 1972). Generally, test results have been similar to those observed by Honig (1969). Thomas, Freeman, Svinicki, Burr, & Lyons (1970) carried out a particularly significant set of experiments using a novel procedure in which at least one of the extradimensional stimuli was compounded with the criterion stimulus during dimensional acquisition. In one study, pigeons first learned to discriminate between two line orientations. *S*+ from that discrimination was then superimposed on a background field of 550 nm during dimensional acquisition. One might expect that the presence of *S*+ would have blocked acquisition of stimulus control by the background color, but it did not. To the contrary, spectral gradients (obtained without the line on the key) were steeper following true- than following pseudo-discrimination training. The same result was obtained for the gradients of line orientation when the roles of the colors and lines were reversed during training.

Thomas et al. (1970) ran an additional pair of studies with a *compounded-cues* design, in which extradimensional and dimensional cues were combined throughout the entire training procedure. Two line orientations were superimposed on a green background during true- or pseudo-discrimination training, and this was followed by a spectral generalization test. True-discrimination training again yielded a steeper spectral gradient, the opposite from what would be expected if cues that are better predictors of reinforcement command more attention than other, less predictive cues.

"Attentional enhancement" as an interpretation. The outcome of these particular experiments led Thomas (1970) to suggest an interpretation of extradimensional training effects that differs markedly from a version based on selective attention. His central supposition is that true-discrimination training generally enhances, rather than diminishes, attention to other stimuli in the environment. When such stimuli are introduced together with or after extradimensional training, they gain more control if such training involves true rather than pseudo-discrimination. This attentional enhancement, considered by Thomas to be an active process, can explain other findings reviewed in this section (e.g., the results from interdimensional procedures). But it will not explain the reciprocal interaction between the light and tone cues in the Miles and Jenkins (1973) study. Furthermore, it stands in direct opposition to the findings of Wagner et al. (1968), whose study was a virtual counterpart to that of Thomas et al. (1970). The line cue in the Thomas et al. experiments served the same function as the light cue in the Wagner et al. research. The latter was presented with either of two tones, the former with either of two colors. One simple explanation for the discrepancy is to argue that stimuli from the same modality (such as visual stimuli like lines and colors) are subject to attentional enhancement, whereas stimuli from different modalities are not. One experiment by Thomas et al. (1970), however, speaks against this view. Thomas et al. trained pigeons on either a true or a pseudo-discrimination between different floor tilts of the experimental chamber while they responded to a 555-nm stimulus. Spectral generalization gradients were clearly steeper following true-discrimination training.

Selective attention and "stimulus X." Mackintosh has offered an alternative explanation of extradimensional effects such as these, as well as those falling under the category of selective attention (Mackintosh, 1974, 1977; Turner & Mackintosh, 1972). Mackintosh's approach is based on an analysis originally proposed by Wagner (1969), and it is most readily applied to those studies in which extradimensional training and dimensional acquisition are carried out in separate phases, or at least on separate trials. Assume that certain unspecified but constant cues are present in both phases of such a study. Let us call such a cue "stimulus X." According to Wagner, stimulus X would fail to gain much control during true-discrimination training because *S*+ and *S*− are more predictive of reinforcement. Stimulus X is still present during subsequent dimensional acquisition but, because it has acquired little control, the dimensional stimulus is able to gain the majority of control in competition with it. This is manifested through a steep generalization gradient. In pseudo-discrimination training, the extradimensional stimuli are no more predictive of reinforcement than stimulus X, so the latter acquires considerable control over responding. This control then limits that acquired by the dimensional value during dimensional acquisition, making the slope of the generalization gradient relatively shallow.

Wagner (1969) has provided some empirical support for this view. Rabbits were given classical eyeblink conditioning with "stimulus X" made explicit in the form of a vibrator applied to the animal's chest on all training trials. A 3100-Hz tone served as the CS in dimensional acquisition, while a steady and a flashing light were the extradimensional CSs. For the true-discrimination group, only one of the lights was paired with the UCS. For the pseudo-discrimination group, both lights were paired with the UCS. Three tonal frequencies were presented during subsequent testing, half the time in combination with the vibrotactile stimulus, the other half without it. Figure 14 shows that the auditory generalization gradients were steeper following true discrimination training, in agreement with the findings of Thomas et al. (1970). Of particular interest is the effect of omitting the vibrotactile cue during testing. First, the tonal gradients were steeper, which is not surprising in view of the masking phenomenon discussed earlier. Second, the overall level of responding was reduced. This is also to be expected since the vibrotactile stimulus was a constant cue during training. However, the reduction was less for the true- than for the pseudo-discrimination group, suggesting that "stimulus X" had acquired less control during the former than during the latter discrimination.

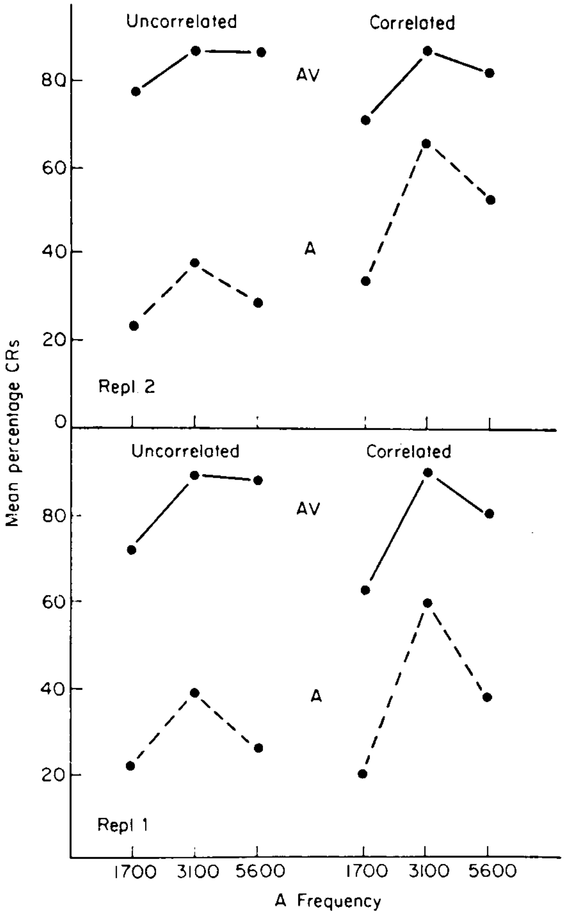

Fig. 14. Mean percentage of eyeblink responses from rabbits during generalization tests which followed conditioning with a compound 3100-Hz tone and vibrotactile cue as CS. For the Correlated group, reinforced trials were signaled by differential visual stimuli. For the Uncorrelated group, visual stimuli were presented randomly on reinforced and nonreinforced trials. The effects of omitting the vibrotactile cue on test trials are particularly evident. See text for further description. (From Wagner, 1969).

Turner and Mackintosh (1972) have suggested a particular source for "stimulus X" in research where it has not been made explicit; namely, the repetitive characteristics of free-operant responding. The studies by Wagner et al. (1968), on which much of the argument in favor of selective attention is based, used discrete-trial methods which preclude repeated responding and thus eliminate that potential source of competing cues. On the other hand, the experiments by Honig (1969) and by Thomas et al. (1970) both involved extended trials with repeated responding.

The easiest way to test whether or not stimuli associated with repetitive responding serve as background cues is simply to remove them. Gray and Mackintosh (1973) did just that by using a discrete-trial procedure where trials ended after a single response (if any was made). A vertical line was presented on each trial as the dimensional training value, while a tone and a white noise served as the extradimensional stimuli. Generalization tests were then run with different line orientations and with the auditory stimuli absent (in order to avoid masking effects). These tests showed that pigeons receiving pseudo-discrimination training responded more to the line and provided much steeper gradients than true-discrimination subjects. This finding is, of course, the opposite of that predicted by attentional enhancement but in accord with the analysis by Wagner (1969).

It is unfortunate that auditory stimuli were used as cues in the Gray and Mackintosh (1973) study, when background colors on the key would have provided more comparability to the Thomas et al. (1970) experiment. It also differed from the usual extended-trial operant research in many ways other than the opportunity for repeated responding (e.g., trials were brief and widely spaced). Hall and Honig (1974), however, carried out a related experiment in which the typical features of operant procedures were maintained, and responding to the key served as an independent variable. Two groups of pigeons served as a response-contingent pair. They were given either true- or pseudo-discrimination training in the presence of different ceiling lights, during which time they pecked at a white response key. Two other groups, the noncontingent pair, were given free reinforcers under similar circumstances but with no light on the response key. After seven discrimination training sessions, all birds were taught to peck at three vertical lines on a black background (the dimensional training value) with the ceiling lights turned off. Generalization tests with different line orientations were then conducted.

The test results, shown in Figure 15, are quite clear. The response-contingent groups provided gradients that differed in accordance with the usual effects of true- and pseudo-discrimination training. Clearly, the pseudo-discrimination gradient was both higher and flatter than the true discrimination gradient. For the noncontingent groups, no such difference was observed; their gradients were almost identical. This experiment, then, strongly suggests that repetitive responding is necessary to obtain attentional enhancement of stimulus control. Such responding provides stimuli common to both extradimensional training and dimensional acquisition, and may thus compete with visual (or other) training stimuli. Removal of the response-contingent cues attenuates the effect.

A complete study of the effects of response-produced cues requires that extradimensional training be separate from dimensional acquisition. Responding to the key, however, must occur at least during acquisition and testing, since it is used to measure stimulus control. Consequently, it would be difficult or impossible to carry out a compounded-cues study to examine the role of repetitive responding in attentional enhancement (i.e., to test the notion that responding serves as "stimulus X" in the very paradigm for which the explanation offered by Mackintosh is most in doubt). Although response-produced cues may determine whether or not attentional enhancement will occur in many extradimensional training experiments, their role in compounded-cues studies remains to be determined. Enhancement in this paradigm may be based upon entirely different processes.

Further support for attentional enhancement. An experiment by Robles, Newlin, and Thomas (1980) provides some support for this last statement. Their design is shown in panel A of Figure 16. In the first stage of training, two separate groups of pigeons learned either a true- or pseudo-discrimination between vertical and horizontal white lines presented on a dark key. These groups are referred to as the *S+ TD* and the *S+ PD* groups, respectively. A second pair of groups was treated similarly, but their stimuli consisted of the illumination of the top and bottom halves of the response key with white light. These are called the *O* (for orthogonal) *TD* and the *O PD* groups. Presumably, preliminary training in both true-discrimination groups should neutralize any incidental cues (e.g., those arising from repetitive key pecking) to the same degree, whereas training in the pseudo-discrimination groups should leave such cues free to gain control

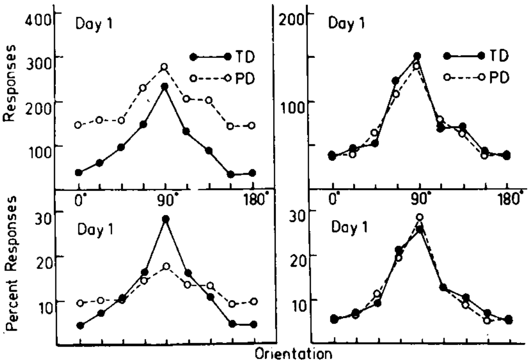

Fig. 15. Absolute and relative line-orientation generalization gradients following either true-discrimination (TD) or pseudo-discrimination (PD) training with overhead colors. Data in the left panels were obtained from pigeons that pecked at a key during initial training with colors. Data in the right panels are from pigeons that received response-independent reinforcement during training. (From Hall & Honig, 1974).

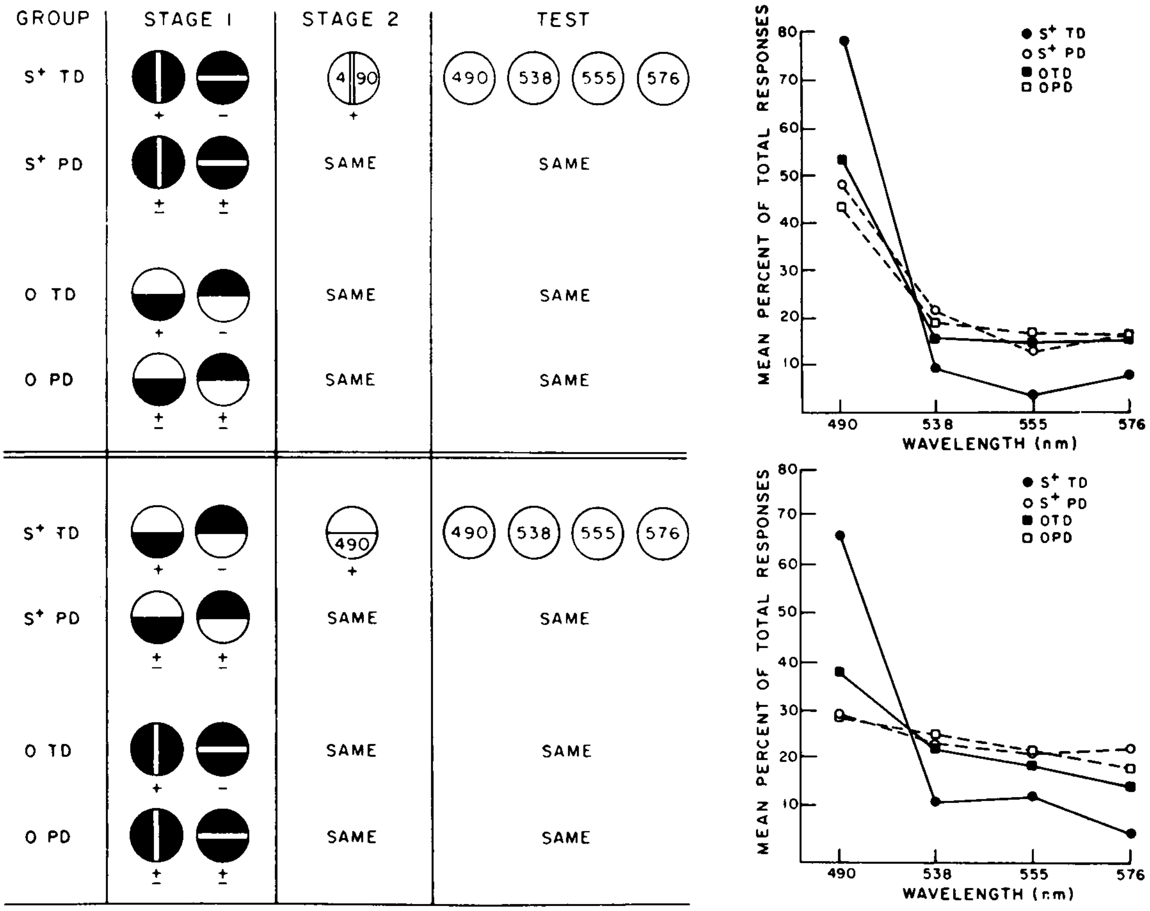

Fig. 16. Relative wavelength generalization gradients obtained after true- and pseudo-discrimination training employing either horizontal and vertical lines or white half-key illumination. The *S*+ line was maintained on the key in dimensional acquisition in Experiment 1. The design is shown in the upper half of Panel A and the results in Panel B. The *S*+ half-field was maintained on the key in dimensional acquisition in Experiment 2. The design is shown in the lower half of Panel A and the results in Panel C. (From Robles, Newlin & Thomas, 1980).

over responding. In the second stage of training, all four groups learned to respond to a 490-nm light upon which the vertical white line was superimposed (dimensional acquisition). Control by wavelength was then assessed by presenting four spectral values (490, 538, 555, and 576 nm) without the vertical line during generalization tests in extinction.

The data from this experiment are shown in panel B of Figure 16. Comparison of the *S*+ *TD* and *S*+ *PD* groups shows a clear enhancement effect: the spectral gradient was steeper after true discrimination than after pseudo-discrimination. A corresponding difference can be seen in the gradients obtained from the *O TD* and the *O PD* groups, although the enhancement effect is much smaller. These results were also obtained in a second study where the roles of the lines and half-fields were reversed. The design of this latter study can be seen in the bottom half of panel A, and the data in panel C.

At first glance, the test results shown in Figure 16 are not surprising. True-discrimination training along one stimulus dimension produces sharper gradients of generalization along an orthogonal dimension than does pseudo-discrimination training. This enhancement effect supposedly occurs because incidental cues in the true-discrimination condition have been neutralized to the extent that they are unable to block and/or mask the control acquired by other stimuli (i.e., wavelength) over the subsequent course of training and testing. The finding of major importance, however, is that the *S*+ *TD* gradient is steeper than the *O TD* gradient. A selective attention model of attentional enhancement predicts that the spectral gradient for the *S+ TD* group should be shallower than for the *O TD* group, because the continued presence of the positive line orientation during dimensional acquisition should block control by wavelength. Blocking is expected on the grounds that the line should have acquired considerable control over responding in Stage 1 at the expense of irrelevant incidental stimuli.

One might reasonably argue, however, that the relatively shallow gradients in the *O TD* groups were artifactual, because color was overshadowed by the novel stimulus in Stage 2 (i.e., the vertical line following the preliminary half-field discrimination, and vice versa). Unfortunately, this argument cannot be resolved on the basis of the available data. Appropriate control groups trained with color alone in Stage 2 are needed. Nonetheless, even if overshadowing were present in the *O TD* groups, it seems unlikely that it would have been strong enough to outweigh any blocking effects present in the *S+ TD* groups. At best, a selective attention model does not easily account for the observed data.

On balance, the Robles et al. experiment suggests that the presence of an *S+* from a prior discrimination can actively enhance attention to a new training value compounded with it. In light of the prior discussion, this raises the possibility that attentional enhancement may be obtained through two different mechanisms. The first mechanism, selective attention, operates when the relevant cues in discrimination training do not overlap with those in dimensional acquisition [as in the Hall and Honig (1974) study]. True discrimination serves to neutralize any incidental cues common to both phases of training. The second mechanism, active enhancement, operates when a cue from discrimination training is carried over into dimensional acquisition. This potentiates the processing of the compounded dimensional stimulus and, consequently, overrides the neutralization of common but irrelevant cues. Although this interpretation is post hoc and far from parsimonious, it may prove to be less cumbersome than any single-process explanation designed to account for the diverse results discussed in this section.

Clearly, further analysis is required. We might reasonably suppose that active enhancement takes place only with stimuli compounded at the same location, such as on the pigeon's pecking key. On the other hand, if two auditory frequencies or two overhead lights were involved in initial true- or pseudodiscrimination training, then the presence of the positive value during dimensional acquisition might detract from (rather than enhance) control by a stimulus on the response key. Indeed, much of the evidence favoring selective attention is derived from studies in which stimuli are presented in different modalities or at least in different locations (Hall & Honig, 1974; Jenkins & Harrison, 1960; Miles & Jenkins, 1973; Wagner et al., 1968; Honig, Note 1). Conversely, studies that favor the notion of active enhancement are similar to the one described by Robles et al. (1980), in that the stimuli from the training and test dimensions are presented in the same place (Honig, 1969; Thomas et al., 1970). The one striking exception, of course, is the experiment reported by Thomas et al. (1970), in which pigeons taught to discriminate between different floor tilts provided steeper spectral generalization gradients than subjects exposed only to a single floor tilt. Obviously, the issues at hand await final resolution.

CONCLUSION

In this review, we have tried to provide a summary of the methods, problems, and findings in the area of stimulus generalization over the past 25 years. We have not advanced a particular theoretical position, nor can we provide a single "take-home" message from that part of the generalization literature with which we have been concerned. The success of the generalization technique has been demonstrated by its application to a wide variety of topics and issues in conditioning and learning. At this point, an overview of the course of developments within the area should provide appropriate closure.

Perhaps of foremost importance, the study of stimulus generalization nicely exemplifies the interplay between methodological advances and theoretical concerns. The original research by Guttman and Kalish was groundbreaking primarily as a methodological contribution. Intermittent reinforcement in training, and the repeated presentation of test stimuli in randomly permuted orders during extinction, provided a method for capturing the gradient even for individual subjects. Yet, the article was written in the context of an historical concern with generalization gradients as theoretical entities. In fact, it began with a sentence pointing out that all possible forms of the generalization gradient had been proposed: concave, linear, and convex. Hull, Spence, and others had used theoretical gradients of these forms in their explanatory systems.

Normally, one would expect that once the gradient became visible, the question of its true form would be clarified. The effect, however, was quite different. The Guttman and Kalish study and those prompted by it clearly showed that the belief in a single gradient form was misleading: it could assume many different forms depending upon the training conditions, stimulus dimension, and testing procedure. Thus, it was pointless to search for an empirical gradient which reflected some underlying, unitary theoretical structure. Instead, the direction of research was guided towards the variables which determine the height, slope, and form of the generalization gradient.

The information obtained from these investigations soon permitted a return to topics and issues which had long remained in the realm of theoretical speculation. The steady-state testing procedure developed by D. Blough is a good example of this point. By using discrete trials and infrequent reinforcement at particular test values, reliable and replicable data concerning generalization within a restricted test range could be obtained. Such would have been impossible with the usual extended-trials extinction procedure. The steady-state method was then used by P. Blough to successfully address the theoretical question posed by Guttman and Kalish in their introduction, namely the relationship between gradient slope and discriminability along the test dimension.

Inhibitory generalization provides still another example of this type of interplay between empirical and conceptual work. Inhibition was also an active issue at the time that Guttman and Kalish conducted their work. Spence had proposed that certain phenomena in discrimination learning resulted from an interaction of theoretical inhibitory gradients with excitatory ones. Hanson initially tried to provide a glimpse of these inhibitory gradients in his study on the peak shift. Unfortunately, his attempted derivation from empirically obtained postdiscrimination and single-stimulus gradients was not particularly satisfying. The subsequent development of orthogonal training procedures, however, finally robbed inhibitory generalization of its mystery. Negative gradients could readily be observed if $S-$ simply varied along a dimension other than that on which $S+$ was located. Attempts were then made, like that by Hearst, to derive the peak shift from independently obtained empirical gradients rather than from theoretical ones. Furthermore, interesting and fruitful discussions arose concerning the fact that negative gradients were often shallower than their positive counterparts. A certain degree of clarification regarding the measurement and operation of inhibitory stimulus control soon came about through modifications of the extinction test procedure (e.g., the resistance-to-reinforcement test) and through the development of rate-independent measures of stimulus control (e.g., stimulus duration).

The problem of attentional factors, unlike some of the others, was not of concern during the early work on generalization. This theoretical issue arose only after it became clear that some stimulus dimensions, such as auditory frequency, yielded flatter gradients than others. The early explanation of this effect was largely comprised of the suggestion that responding to one set of stimuli is controlled in part by other stimuli that are not as readily manipulable. Since we have reviewed in detail the work on this problem, it should be sufficient to point out that once again the development of new methods (in this case, inter- and extradimensional training procedures) shed considerable light on how various stimuli could simultaneously enter into the control over behavior. In particular, research on attentional mechanisms gave some theoretical status to the notion that repetitive responding can be a potent discriminative cue.

Active research on operant stimulus generalization has declined somewhat in recent years. Part of the decline is probably attributable to the fact that certain problems and issues have been thoroughly explored. A more important reason, however, is that the principal dependent variable, rate of responding, has reached its apparent limit of usefulness in the field. This measure simply does not allow an effective separation of those variables which determine it, a separation necessary for more detailed experimental analyses of behavior. This is all a bit ironic, considering that rate of responding was introduced by Guttman and Kalish as their principal methodological innovation. Nonetheless, our own behavior as scientists must be sensitive to the demands for finer analyses and to the wealth of information already obtained through many years of considerable research effort. The methods and measures involved in the study of stimulus control have and will continue to change. The generalization technique has provided us with a solid groundwork on which to build.

REFERENCE NOTE

1. Honig, W. K. Selective attention, attentional enhancement, and stimulus control. Paper read at the 19th International Congress of Psychology, London, England, 1969.

REFERENCES

Baker, W. M., & Holland, M. K. Spectral generalization gradients from pigeons as a function of response latency. *Psychonomic Science,* 1968, **10,** 3-4.

Beale, I. L., & Winton, A. S. W. Inhibitory stimulus control in concurrent schedules. *Journal of the Experimental Analysis of Behavior,* 1970, **14,** 133-137.

Bitterman, M. E. Generalization. In M. E. Bitterman, V. M. LoLordo, J. B. Overmier, & M. E. Rashotte. *Animal learning: Survey and analysis.* New York: Plenum Press, 1979.

Bloomfield, T. M. A peak shift on a line-tilt continuum. *Journal of the Experimental Analysis of Behavior,* 1967, **10,** 361-366.

Blough, D. S. Spectral sensitivity in the pigeon. *Journal of the Optical Society of America,* 1957, **47,** 827-833.

Blough, D. S. The shape of some wavelength generalization gradients. *Journal of the Experimental Analysis of Behavior,* 1961, **4,** 31-40.

Blough, D. S. Interresponse time as a function of continuous variables: A new method and some data. *Journal of the Experimental Analysis of Behavior,* 1963, **6,** 237-246.

Blough, D. S. Definition and measurement in generalization research. In D. I. Mostofsky (Ed.), *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Blough, D. S. Stimulus generalization as signal detection in pigeons. *Science,* 1967, **158,** 940-941.

Blough, D. S. Generalization gradient shape and summation in steady-state tests. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 91-104. (a)

Blough, D. S. Attention shifts in a maintained discrimination. *Science,* 1969, **166,** 125-126. (b)

Blough, D. S. Recognition by the pigeon of stimuli varying in two dimensions. *Journal of the Experimental Analysis of Behavior,* 1972, **18,** 345-367.

Blough, D. S. Two-way generalization peak shift after two-key training in the pigeon. *Animal Learning and Behavior,* 1973, **1,** 171-174.

Blough, D. S. Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology: Animal Behavior Processes,* 1975, **1,** 3-21.

Blough, D. S. Reaction times of pigeons on a wavelength discrimination task. *Journal of the Experimental Analysis of Behavior,* 1978, **30,** 163-167.

Blough, P. M. Wavelength generalization and discrimination in the pigeon. *Perception and Psychophysics,* 1972, **12,** 342-348.

Boneau, C. A., & Cole, J. L. Decision theory, the pigeon, and the psychophysical function. *Psychological Review,* 1967, **74,** 123-135.

Boneau, C. A., & Honig, W. K. Opposed generalization gradients based upon conditional discrimination training. *Journal of Experimental Psychology,* 1964, **68,** 89-93.

Boring, E. G., Langfeld, H. S., and Weld, H. P. *Foundations of psychology.* New York: Wiley, 1948.

Bush, R. R., & Mosteller, F. A mathematical model for simple learning. *Psychological Review,* 1951, **58,** 313-323.

Butter, C. M. Stimulus generalization along one and two dimensions in pigeons. *Journal of Experimental Psychology,* 1963, **65,** 339-346.

Chase, S., & Heinemann, E. G. Choices based on redundant information: An analysis of two-dimensional stimulus control. *Journal of Experimental Psychology,* 1972, **92,** 161-175.

Deutsch, J. A. Discrimination learning and inhibition. *Science,* 1967, **156,** 988.

Engen, T. Psychophysics. 1. Discrimination and detection. In J. W. Kling & L. A. Riggs (Eds.), *Woodworth and Schlosberg's experimental psychology.* New York: Holt, Rinehart, & Winston, 1971.

Farthing, G. W. Overshadowing in the discrimination of successive compound stimuli. *Psychonomic Science,* 1972, **28,** 29-32.

Friedman, H. Wave-length generalization as a function of spacing of test stimuli. *Journal of Experimental Psychology,* 1963, **65,** 334-338.

Friedman, H., & Guttman, N. Further analysis of the various effects of discrimination training on stimulus generalization gradients. In D. I. Mostofsky (Ed.), *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Ganz, L. Hue generalization and hue discriminability in *Macaca Mulatta. Journal of Experimental Psychology,* 1962, **64,** 142-150.

Ganz, L., & Riesen, A. H. Stimulus generalization to hue in the dark-reared macaque. *Journal of Comparative and Physiological Psychology,* 1962, **55,** 92-99.

Gray, V. A. Stimulus control of differential-reinforcement-of-low-rate responding. *Journal of the Experimental Analysis of Behavior,* 1976, **25,** 199-207.

Gray, V. A., & Mackintosh, N. J. Control by an irrelevant stimulus in discrete-trial discrimination learning by pigeons. *Bulletin of the Psychonomic Society,* 1973, **1,** 193-195.

Guttman, N. Generalization gradients around stimuli associated with different reinforcement schedules. *Journal of Experimental Psychology,* 1959, **58,** 335-340.

Guttman, N. Laws of behavior and facts of perception. In S. Koch (Ed.), *Psychology: A study of a science* (Vol. 5). New York: McGraw-Hill, 1963.

Guttman, N. Effects of discrimination formation on generalization measured from the positive-rate baseline. In D. I. Mostofsky (Ed.), *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Guttman, N., & Kalish, H. I. Discriminability and stimulus generalization. *Journal of Experimental Psychology,* 1956, **51,** 79-88.

Guttman, N., & Kalish, H. I. Experiments in discrimination. *Scientific American,* 1958, **198,** No. 1, 77-82.

Hall, G., & Honig, W. K. Stimulus control after extradimensional training in pigeons: A comparison of response contingent and noncontingent training procedures. *Journal of Comparative and Physiological Psychology,* 1974, **87,** 945-952.

Hamilton, W. F., & Coleman, T. B. Trichromatic vision in the pigeon as illustrated by the spectral hue discrimination curve. *Journal of Comparative Psychology,* 1933, **15,** 183-191.

Hanson, H. M. Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology,* 1959, **58,** 321-334.

Hanson, H. M. Stimulus generalization following three-stimulus discrimination training. *Journal of Comparative and Physiological Psychology,* 1961, **54,** 181-185.

Hearst, E. Discrimination learning as the summation of excitation and inhibition. *Science,* 1968, **162,** 1303-1306.

Hearst, E. Excitation, inhibition, and discrimination learning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning.* Halifax: Dalhousie University Press, 1969.

Hearst, E. Some persistent problems in the analysis of conditioned inhibition. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning.* New York: Academic Press, 1972.

Hearst, E., Besley, S., & Farthing, G. W. Inhibition and the stimulus control of operant behavior. *Journal of the Experimental Analysis of Behavior,* 1970, **14,** 373-409.

Hearst, E., Koresko, M. B., & Poppen, R. Stimulus generalization and the response-reinforcement contingency. *Journal of the Experimental Analysis of Behavior,* 1964, **7,** 369-380.

Heinemann, E. G., Avin, E., Sullivan, M. A., & Chase, S. Analysis of stimulus generalization with a psychophysical method. *Journal of Experimental Psychology,* 1969, **80,** 215-224.

Heinemann, E. G., & Chase, S. Conditional stimulus control. *Journal of Experimental Psychology,* 1970, **84,** 187-197.

Heinemann, E. G., & Chase, S. Stimulus generalization. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (Vol. 2). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1975.

Hiss, R. H., & Thomas, D. R. Stimulus generalization as a function of testing procedure and response measure. *Journal of Experimental Psychology,* 1963, **65,** 587-592.

Hoffman, H. S., & Fleshler, M. Stimulus factors in aversive controls: The generalization of conditioned suppression. *Journal of the Experimental Analysis of Behavior,* 1961, **4,** 371-378.

Holland, M. K., & Baker, W. M. Spectral generalization testing with pigeons using brief discrete trials. *Psychonomic Science,* 1968, **10,** 1-2.

Honig, W. K. Generalization of extinction on the spectral continuum. *Psychological Record,* 1961, **11,** 269-278.

Honig, W. K. Prediction of preference, transposition, and transposition-reversal from the generalization gradient. *Journal of Experimental Psychology,* 1962, **64,** 239-248.

Honig, W. K. Attentional factors governing the slope of the generalization gradient. In R. M. Gilbert and N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969.

Honig, W. K. Attention and the modulation of stimulus control. In D. Mostofsky (Ed.), *Attention: Contemporary theory and analysis.* New York: Appleton-Century-Crofts, 1970.

Honig, W. K. Effects of extradimensional discrimination training upon previously acquired stimulus control. *Learning and Motivation,* 1974, **5,** 1-15.

Honig, W. K., & Beale, I. L. Stimulus duration as a measure of stimulus generalization. *Journal of the Experimental Analysis of Behavior,* 1976, **25,** 209-217.

Honig, W. K., Beale, I., Seraganian, P., Lander, D., & Muir, D. Stimulus and response reduction: Two aspects of inhibitory control in learning. In R. A. Boakes & M. S. Halliday (Eds.), *Inhibition and learning.* New York: Academic Press, 1972.

Honig, W. K., Boneau, C. A., Burstein, K. R., & Pennypacker, H. S. Positive and negative generalization gradients obtained after equivalent training conditions. *Journal of Comparative and Physiological Psychology,* 1963, **56,** 111-116.

Honig, W. K., & Lindsay, H. Transfer of a response controlling stimulus duration across discrimination problems. *Learning and Motivation,* 1975, **6,** 157-178.

Honig, W. K., Thomas, D. R., & Guttman, N. Differential effects of continuous extinction and discrimination training on the generalization gradient. *Journal of Experimental Psychology,* 1959, **58,** 145-152.

Hull, C. L. *Principles of behavior.* New York: Appleton-Century-Crofts, 1943.

Hull, C. L. The problem of primary stimulus generalization. *Psychological Review,* 1947, **54,** 120-134.

Jenkins, H. M. Generalization gradients and the concept of inhibition. In D. I. Mostofsky (Ed.), *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Jenkins, H. M., & Harrison, R. H. Effect of discrimination training on auditory generalization. *Journal of Experimental Psychology,* 1960, **59,** 246-253.

Jenkins, H. M., & Harrison, R. H. Generalization gradients of inhibition following auditory discrimination learning. *Journal of the Experimental Analysis of Behavior,* 1962, **5,** 435-441.

Jitsumori, M. Wavelength discrimination function derived from post-discrimination gradients in the pigeon. *Japanese Psychological Research,* 1978, **20,** 18-28.

Johnson, D. F. Determiners of selective stimulus control in the pigeon. *Journal of Comparative and Physiological Psychology,* 1970, **70,** 298-307.

Kalish, H. I. The relationship between discriminability and generalization: A re-evaluation. *Journal of Experimental Psychology,* 1958, **55,** 637-644.

Kalish, H. I. Stimulus generalization. In M. H. Marx (Ed.), *Learning: Processes.* London: Collier-MacMillan, 1969.

Kalish, H. I., & Guttman, N. Stimulus generalization after equal training on two stimuli. *Journal of Experimental Psychology,* 1957, **53,** 139-144.

Kalish, H. I., & Guttman, N. Stimulus generalization after training on three stimuli: A test of the summation hypothesis. *Journal of Experimental Psychology,* 1959, **57,** 268-272.

Kalish, H. I., & Haber, A. Generalization: I. Generalization gradients from single and multiple stlmulus points. II. Generalization of inhibition. *Journal of Experimental Psychology,* 1963, **65,** 176-181.

Kamin, L. J. "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior, 1967: Aversive stimulation.* Miami: University of Miami Press, 1968.

Karpicke, J., & Hearst, E. Inhibitory control and errorless discrimination learning. *Journal of the Experimental Analysis of Behavior,* 1975, **23,** 159-166.

Lashley, K. S., & Wade, M. The Pavlovian theory of generalization. *Psychological Review,* 1946, **53,** 72-87.

Lovejoy, E., & Russell, D. G. Suppression of learning about a hard cue by the presence of an easy cue. *Psychonomic Science,* 1967, **8,** 365-366.

Lyons, J., & Thomas, D. R. Effects of interdimensional training on stimulus generalization: II. Within-subjects design. *Journal of Experimental Psychology,* 1967, **75,** 572-574.

Mackintosh, N. J. *The psychology of animal learning.* London: Academic Press, 1974.

Mackintosh, N. J. Stimulus control: Attentional factors. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior.* Englewood Cliffs, N.J.: Prentice-Hall, 1977.

Mackintosh, N. J., & Honig, W. K. Blocking and attentional enhancement in pigeons. *Journal of Comparative and Physiological Psychology,* 1970, **73,** 78-85.

Malott, M. K. Stimulus control in stimulus-deprived chickens. *Journal of Comparative and Physiological Psychology,* 1968, **66,** 276-282.

Marsh, G. D. Inverse relationship between discriminability and stimulus generalization as a function of number of test stimuli. *Journal of Comparative and Physiological Psychology,* 1967, **64,** 284-289.

Marsh, G. D. Prediction of the peak shift in pigeons from gradients of excitation and inhibition. *Journal of Comparative and Physiological Psychology,* 1972, **81,** 262-266.

Miles, C. G., & Jenkins, H. M. Overshadowing in operant conditioning as a function of discriminability. *Learning and Motivation,* 1973, **4,** 11-27.

Miller, N. E. Experimental studies of conflict. In J. McV. Hunt (Ed.), *Personality and the behavior disorders.* New York: Ronald Press, 1944.

Morgan, M. J. Book review of *Animal discrimination learning,* R. M. Gilbert & N. S. Sutherland (Eds.). *Quarterly Journal of Experimental Psychology,* 1969, **21,** 291-292.

Mostofsky, D. I. (Ed.). *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Mountjoy, P. T., & Malott, M. K. Wavelength generalization curves for chickens reared in restricted portions of the spectrum. *Psychological Record,* 1968, **18,** 575-583.

Newman, F. L., & Baron, M. R. Stimulus generalization along the dimension of angularity. *Journal of Comparative and Physiological Psychology,* 1965, **60,** 59-63.

Pavlov, I. P. *Conditioned reflexes.* Oxford: Oxford University Press, 1927.

Peterson, N. Effect of monochromatic rearing on the control of responding by wavelength. *Science,* 1962, **136,** 774-775.

Razran, G. Stimulus generalization of conditioned responses. *Psychological Bulletin,* 1949, **46,** 337-365.

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

Reynolds, G. S. Behavioral contrast. *Journal of the Experimental Analysis of Behavior,* 1961, **4,** 57-71. (a)

Reynolds, G. S. Attention in the pigeon. *Journal of the Experimental Analysis of Behavior,* 1961, **4,** 203-208. (b)

Riley, D. A. *Discrimination learning.* Boston: Allyn & Bacon, 1968.

Rilling, M. Stimulus control and inhibitory processes. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior.* Englewood Cliffs, N.J.: Prentice-Hall, 1977.

Rilling, M., Caplan, H. J., Howard, R. C., & Brown, C. H. Inhibitory stimulus control following errorless discrimination learning. *Journal of the Experimental Analysis of Behavior,* 1975, **24,** 121-133.

Robles, L., Newlin, R. J., & Thomas, D. R. Potentiation of attentional enhancement in the pigeon produced by a "blocking" stimulus. *Animal Learning and Behavior,* 1980, **8,** 31-36.

Rosen, A. P., & Terrace, H. S. On the minimal conditions for the development of a peak-shift and inhibitory stimulus control. *Journal of the Experimental Analysis of Behavior,* 1975, **23,** 385-414.

Rudolph, R. L., & Honig, W. K. Effects of monochromatic rearing on spectral discrimination learning and the peak shift in chicks. *Journal of the Experimental Analysis of Behavior,* 1972, **17,** 107-111.

Rudolph, R. L., Honig, W. K., & Gerry, J. E. Effects of monochromatic rearing on the acquisition of stimulus control. *Journal of Comparative and Physiological Psychology,* 1969, **67,** 50-57.

Rudolph, R. L., & Van Houten, R. Auditory stimulus control in pigeons: Jenkins & Harrison (1960) revisited. *Journal of the Experimental Analysis of Behavior,* 1977, **27,** 327-330.

Schlosberg, H., & Solomon, R. L. Latency of response in a choice discrimination. *Journal of Experimental Psychology,* 1943, **33,** 22-39.

Schwartzbaum, J. S., & Kellicutt, M. H. Inverted generalization gradients about a nonreinforced stimulus. *Psychological Reports,* 1962, **11,** 791-792.

Shepard, R. N. Approximation to uniform gradients of generalization by monotone transformations of scale. In D. I. Mostofsky (Ed.), *Stimulus generalization.* Stanford: Stanford University Press, 1965.

Skinner, B. F. Are theories of learning necessary? *Psychological Review,* 1950, **57,** 193-216.

Skinner, B. F. *Science and human behavior.* New York: Macmillan, 1953.

Spence, K. W. The differential response of animals to stimuli differing within a single dimension. *Psychological Review,* 1937, **44,** 430-444.

Sutherland, N. S., & Mackintosh, N. J. *Mechanisms of animal discrimination learning.* New York: Academic Press, 1971.

Switalski, R. W., Lyons, J., & Thomas, D. R. Effects of interdimensional training on stimulus generalization. *Journal of Experimental Psychology,* 1966, **72,** 661-666.

Terrace, H. S. Discrimination learning with and without "errors." *Journal of the Experimental Analysis of Behavior,* 1963, **6,** 1-27. (a)

Terrace, H. S. Errorless transfer of a discrimination across two continua. *Journal of the Experimental Analysis of Behavior,* 1963, **6,** 223-232. (b)

Terrace, H. S. Wavelength generalization after discrimination learning with and without errors. *Science,* 1964, **144,** 78-80.

Terrace, H. S. Stimulus control. In W. K. Honig (Ed.), *Operant behavior.* New York: Appleton-Century-Crofts, 1966. (a)

Terrace, H. S. Discrimination learning and inhibition. *Science,* 1966, **154,** 1677-1680. (b)

Terrace, H. S. Behavioral contrast and peak shift: Effects of extended discrimination training. *Journal of the Experimental Analysis of Behavior,* 1966, **9,** 613-617. (c)

Terrace, H. S. Discrimination learning, the peak shift, and behavioral contrast. *Journal of the Experimental Analysis of Behavior,* 1968, **11,** 727-741.

Terrace, H. S. Evidence for the innate basis of the hue dimension in the duckling. *Journal of the Experimental Analysis of Behavior,* 1975, **24,** 79-87.

Thomas, D. R. Stimulus selection, attention, and related matters. In J. H. Reynierse (Ed.), *Current issues in animal learning.* Lincoln: University of Nebraska Press, 1970.

Thomas, D. R., Freeman, F., Svinicki, J. G., Burr, D. E. S., & Lyons, J. Effects of extradimensional training on stimulus generalization. *Journal of Experimental Psychology,* Monograph, 1970, **83,** 1-21.

Thomas, D. R., & King, R. A. Stimulus generalization as a function of level of motivation. *Journal of Experimental Psychology,* 1959, **57,** 323-328.

Thomas, D. R., & Switalski, R. W. Comparison of stimulus generalization following variable-ratio and variable-interval training. *Journal of Experimental Psychology,* 1966, **71,** 236-240.

Tomie, A., Davitt, G. A., & Engberg, L. A. Stimulus generalization of auto-shaped key-pecking following interdimensional and extradimensional training. *Learning and Motivation,* 1976, **7,** 246-253.

Tracy, W. K. Wavelength generalization and preference in monochromatically reared ducklings. *Journal of the Experimental Analysis of Behavior,* 1970, **13,** 163-178.

Turner, C., & Mackintosh, N. J. Stimulus selection and irrelevant stimuli in discrimination learning by pigeons. *Journal of Comparative and Physiological Psychology,* 1972, **78,** 1-9.

vom Saal, W., & Jenkins, H. M. Blocking the development of stimulus control. *Learning and Motivation,* 1970, **1,** 52-64.

Wagner, A. R. Incidental stimuli and discrimination learning. In R. M. Gilbert & N. S. Sutherland (Eds.), *Animal discrimination learning.* London: Academic Press, 1969.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology,* 1968, **76,** 171-180.

Weisman, R. G. Some determinants of inhibitory stimulus control. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 443-450.

Weisman, R. G., & Palmer, J. A. Factors influencing inhibitory stimulus control: Discrimination training and prior non-differential reinforcement. *Journal of the Experimental Analysis of Behavior,* 1969, **12,** 229-237.

Wright, A. A. Psychometric and psychophysical hue discrimination functions for the pigeon. *Vision Research,* 1972, **12,** 1447-1464.

Wright, A. A. A psychometric and psychophysical theory within a framework of response bias. *Psychological Review,* 1974, **81,** 322-347.

Wright, A. A., & Cumming, W. W. Color-naming functions for the pigeon. *Journal of the Experimental Analysis of Behavior,* 1971, **15,** 7-17.

Zeki, S. The representation of colours in the cerebral cortex. *Nature,* 1980, **284,** 412-418.

Zentall, T., Collins, N., & Hearst, E. Generalization gradients around a formerly positive S–. *Psychonomic Science,* 1971, **22,** 257-259.

Received September 26, 1980
Final acceptance June 12, 1981