

GENERALIZATION GRADIENT SHAPE AND SUMMATION IN STEADY-STATE TESTS¹

DONALD S. BLOUGH

BROWN UNIVERSITY

Pigeons' pecks at one or two wavelengths were reinforced intermittently. Random series of adjacent wavelengths appeared without reinforcement. Gradients of responding around the reinforced wavelengths were allowed to stabilize over a number of sessions. The single (one reinforced stimulus) and summation (two reinforced stimuli) gradients were consistent with a statistical decision account of the generalization process.

Generally speaking, if an operant response is reinforced in the presence of a stimulus, the response will later occur with relatively high probability when that stimulus is present. What happens if the response is reinforced in the presence of two or more stimuli? Two cases, limited to two stimuli for simplicity, may be distinguished: (a) Responses are measured to the two stimuli presented singly, and to the two presented together. If joint stimulation produces a greater response than either stimulus alone, "summation" is said to have occurred. (b) Responses are measured to a set of unreinforced stimuli related in some way to the two reinforced stimuli. The response to each of the unreinforced stimuli yields one point on a "generalization gradient". Gradients following single-stimulus reinforcement are compared with those following reinforcement of two stimuli. If the two-stimulus gradient is higher at some point than either single-stimulus gradient, "generalization gradient summation" is said to have occurred at that point.

This paper is about generalization gradient summation, but no implications are intended about the nature of the discriminative processes that may be involved. Unfortunately, terms concerned with stimulus control have not been precisely used, and they often have unintended theoretical connotations. "Gener-

alization" and "discrimination" are notorious examples, since they do not even imply distinct procedures. The confusing consequences are typified in the present studies, which some readers might prefer to call "generalization" experiments, others "discrimination" experiments.

A distinction of value, however, is that between transient and steady-state situations (Sidman, 1960). Conditioning and extinction typify behavioral transients. The generalization method pioneered by Guttman and Kalish (1956) shows stimulus control during the extinction transient. Though the transient data from the Guttman and Kalish method have provided important qualitative information, attempts to state the shape of functions and to account for combination effects such as gradient summation (Kalish and Guttman, 1957, 1959) and generalization peak shift (Hanson, 1959) have been inconclusive. Perhaps this results partly from the fact that the measurements involve complex changes and response components that are difficult to sort out (Blough, 1963).

The complexities introduced by transient testing may be reduced by treating generalization as a psychophysical problem (Boneau and Cole, 1967; Blough, 1967). For this purpose, gradients are obtained repeatedly under the same conditions for many sessions (Pierrel, 1958). The present experiments generated steady-state gradients centered around several reinforced stimuli on a wavelength dimension. A recently developed reinforcement schedule (Blough, 1966) kept responding stable yet relatively unstereotyped throughout the experiment. Both single and summated gradients

¹Dedicated to B. F. Skinner in his sixty-fifth year. This research was supported in part by USPHS grants MH-02456 and MH-08355. Mrs. Patricia Blough contributed to the research by running the subjects and computer programming. Reprints may be obtained from the author, Walter S. Hunter Laboratory of Psychology, Brown University, Providence, Rhode Island 02912.

were collected for each of three subjects; these were regularized somewhat by rescaling the stimuli, and their form was related to possible controlling mechanisms.

METHOD

Subjects

Three White Carneaux pigeons were maintained at approximately 75% of free-feeding weight by supplementary feeding, if necessary, after each experimental session. The birds had all served previously in an experiment on reinforcement schedules (Blough, 1966), and Birds 556 and 812 had a variety of discrimination training prior to that.

Apparatus

The birds worked simultaneously in three standard Grason Stadler pigeon chambers. These chambers were dark except for a stimulus spot projected upon the response key. The key switches closed on application of about 10-g force. Each chamber was equipped with a ventilating fan and a loudspeaker that supplied white masking noise.

Stimuli were provided by a Bausch & Lomb 250-mm grating monochromator equipped with a ribbon filament lamp operated at 17 amp ac. Entrance and exit slits were set to provide a half-width dispersion of 6.6 nm. Light of this limited wavelength passed through a set of shutters to focus on the ends of three fiber-optics light guides of $\frac{1}{8}$ in. diameter by 3 ft. length. The terminal ends of the light guides were mounted approximately $\frac{1}{4}$ in. behind the translucent plastic response keys in the pigeon chambers. The stimulus patch supplied by the wires consisted of a fuzzy spot of approximately $\frac{3}{8}$ in. diameter, centered on the circular response key. The luminance of the spot was approximately six foot lamberts at 580 nm. The spot appeared on the key at all times except during reinforcement, during intertrial intervals, and for 0.6 sec after each peck. The key went dark after each peck to provide feedback to the subject for effective responses, and because responses at shorter interresponse times (IRTs) are not generally under stimulus control (Blough, 1963, 1966).

A small general-purpose computer, the LINC, (Clark and Molnar, 1964) sensed closures of the key switches. The LINC stored

response latencies and interresponse times in its magnetic core memory; the data were periodically written onto magnetic tape. During the session, the LINC programmed stimulus presentations, drove a stepping motor to adjust stimulus wavelength, and operated shutters to control presentation intervals. The machine also delivered reinforcements via its output relays, according to programmed instructions outlined below. Following each session, the LINC printed key data tables and graphs on a teletype; subsequently, other data analyses were performed and the results either printed or graphically displayed on the machine's oscilloscope.

Procedure

The experiment ran daily for seven months, each daily session lasting 134 min. The session consisted of a sequence of 30-sec trials, during which the stimulus spot illuminated the response key. Three-second blackout periods separated the trials. On some of the trials reinforcement was available on an intermittent schedule; these will be called "S+" trials. These S+ trials were mixed with unreinforced trials ("test trials") according to a semi-random sequence as follows. Each session began with four S+ trials. Following this warmup, the session was divided into 15 stimulus sequences, presented serially. Each sequence consisted of 12 test trials, each a different wavelength, and four S+ trials, all 16 mixed in random order. It is important to note that if any wavelength appeared in the sequence as an S+, it also appeared once, unreinforced, as a test stimulus. The data presented in this paper are from test trials only.

Reinforcement consisted of 3-sec access to mixed grain. The reinforcement schedule in effect during S+ presentations was a somewhat simplified version of a schedule, described in detail elsewhere (Blough, 1966), called the "reinforcement of least-frequent interresponse times" or "LF" schedule. This schedule reinforced only those responses that terminated IRTs that the bird emitted least often, relative to the distribution of IRTs that would be expected were responses to occur on the average of 1 per sec, but randomly in time. On S+ trials, each peck that met the IRT criterion momentarily in force produced reinforcement. For reinforcement purposes, the latency of a peck from the beginning of a trial counted as

an IRT. After each reinforcement and each trial, the LINC recomputed the IRT next required for reinforcement. Only response data from previous S+ trials entered the computation, and of these, only the most-recent 72 IRTs. For a full explanation of the details of this computation, see Blough (1966). In respects here unspecified, the schedule was as reported in that paper. The birds obtained approximately 60 to 70 reinforcements per session.

The LF schedule produces a pattern of reinforcement and of behavior much like that of a variable-interval schedule. It has the advantage for the present study of maintaining relatively constant response rates across subjects and through long experimental procedures. In previous work (unpublished) similar to that reported here, VI schedules have produced rates that climbed dramatically from session to session, in some birds approaching the high rate characteristically generated by a ratio schedule.

Single and double S+ sessions. As outlined above, each session started with four S+ trials and continued with 15 series of 16 trials each, a series comprising 12 test trials and four S+ trials. For single S+ runs, just one of the test wavelengths appeared on the four S+ trials in each sequence. During these S+ trials, reinforcement was available on the schedule just described.

During most of the experiment, two wavelengths rather than a single one, were selected for reinforcement. During these sessions, each of these two stimuli appeared on two of the four warmup S+ trials, and each appeared on two of the four S+ trials interspersed with the 12 test trials of each sequence. The LF reinforcement schedule was maintained independently for each of the two S+ wavelengths.

Stimulus ranges and S+ placements. Stimulus wavelengths over a 44-nm range from 558 nm to 602 nm were presented during the experiment. This range was selected because of the relatively high energies available from the monochromator, and the high and relatively constant luminance of the stimuli for the pigeon (Blough, 1958) over this part of the spectrum. Sometimes the entire range was used, with the 12 test stimuli spaced 4 nm apart across the spectrum; on other occasions, a narrower part of the spectrum was used, with correspondingly closer spacing of the stimuli. In

general, the wide range was employed with two widely spaced S+s, and the range narrowed for a single S+ or two closely spaced S+s. Each range and combination of S+s appeared repeatedly until the data collected from day to day appeared to be stable. Table 1 lists the stimulus conditions and the number of days spent on each.

Table 1
Stimulus Conditions Employed in the Experiment

<i>S+ in nm</i>	<i>Test stimuli, nm</i>	<i>Number of sessions</i>
570 (pretraining)	570 only	17
570	560-582, 2-nm steps	43*
570	558-602, 4-nm steps	13
570, 590	same	27*
570, 586	same	20*
574, 586	same	15*
578, 586	562-598 in 4-nm steps, plus 580, 584	15*
580, 584	same	4
same	566-594 in 4-nm steps, plus 576, 580, 584, 588	12*
582	same	2
same	570-592 in 2-nm steps	8*
578, 586	same	2
574, 586	558-602 in 4-nm steps	1
570, 590	same	26*
590	578-600 in 2-nm steps	24*
582	576-587 in 1-nm steps	31*

*Last six sessions analyzed and shown in Figures.

As Table 1 shows, single S+ gradients were obtained around three wavelengths, 570, 582, and 590 nm. Double S+ gradients were obtained around combinations of these and intermediate wavelengths. In a number of cases (see Table 1), a few transition sessions intervened between prolonged runs on a given condition. These helped to prevent the subjects' behavior from being disrupted, particularly when marked shifts in S+ wavelength were in progress.

RESULTS

Single Gradients

Each bird produced single S+ gradients around three wavelengths, 570 nm, 582 nm, and 590 nm, with test stimuli spaced at 2-nm intervals. The gradient around 582 nm was replicated with 1-nm spacing in two birds. Figure 1 shows the results of all these single S+ conditions. The data are means over the




Fig. 1. Gradients around single reinforced stimuli located at three points on the wavelength continuum. Each function represents the mean of six sessions. Test stimuli were spaced at 2-nm intervals except for replications (filled circles) with stimuli 1-nm apart.

last six days of the given procedure. These and subsequent data include only responses from unreinforced test presentations, and also omit responses to the first of the 15 daily stimulus series. In all graphs, "response rate" means responses per minute when the stimulus was on; this quantity includes a correction for the 0.6-sec off period after each peck, since such off periods were sometimes extended by rapid multiple pecks.

Figure 1 shows that the birds gave rather regular gradients of similar form. Individual differences across birds are largely consistent differences in the width of the functions, those of Bird 812 being relatively wide and those of Bird 556 relatively narrow. Across wavelengths, a consistent effect is seen with all the birds: the curves tend to be steeper to the right of the reinforced wavelength than to the left, and they are narrower at the right of the

figure (around 590 nm) than the left (around 570 nm). Both of these observations are consistent with the idea that equal wavelength steps are not the most appropriate way to space the stimuli. Hence, rescaling of the abscissa was attempted, as described below.

Gradients with two S+s. Considered now are gradients generated by the reinforcement of responses at two stimulus wavelengths, rather than just one. Once again, the data consist of mean response frequencies for each bird over the final six days of a given reinforcement condition. On the basis of these sets of data, together with the single gradient data presented in Fig. 1, a rough rescaling of the stimulus continuum was attempted along the lines suggested by Shepard (1965). Since the single gradients did not overlap extensively, and the bimodal gradients posed a complex analytical problem, an analytical solution that would yield the desired uniform gradient shape was not attempted. Instead, a graphical approximation was performed which resulted essentially in stretching the continuum at the longer wavelength end and compressing it at the shorter, to yield roughly symmetrical curves. The single-stimulus gradients were plotted on cumulative probability paper, and the two limbs around each S+ were fitted with straight lines. The abscissa scale was adjusted to make the slope of these lines as equal as possible. Since there were some apparently systematic differences between birds, the resulting scale is not optimum for any individual bird.

The dual S+ gradients are shown on this new abscissa in Fig. 2. The curves in this figure include one single S+ gradient around 582 nm (leftmost curve), and also include the replication of the widest separation of S+s (570 nm and 590 nm), shown by the dashed curve in the rightmost rank. The implications of these curves will be touched on below. Note the individual differences, with Bird 812 yielding unimodal curves for all conditions (except a suggestion of bimodality in the replication curve) and the other birds showing a regular progression from a bimodal to a unimodal curve.

Components of gradients. The shape of some of the gradients in Fig. 1 and 2, especially the flat ones of Bird 017, prompts inquiry about possible components that may combine to produce these curves. Such inquiry

is appropriate, too, because this study attempts to approach as simple a steady-state situation as possible, eliminating sequential components inherent in the standard generalization testing procedure.

The data were examined for the effects of several possible variables that might be expected to introduce variability in the curves. No trend could be detected in the data across the six days that entered into the means shown in Fig. 1, 2, 3, and 4, nor were consistent within-session effects evident in data over that portion of each session (the last 14 stimulus series) presented here. However, a consistent effect appeared when the results were broken down by time within 30-sec stimulus presentations. Responding when the stimuli first appeared was less well-controlled (flatter gradients) than was responding after the stimulus had been present for some seconds. This effect can be seen in Fig. 3, where responses occurring in successive thirds of each 30-sec trial are segregated, and functions constructed for each time period. This sample, consisting of gradients around 590 nm, shows the first 10-sec gradients higher in two cases, but in all cases proportionally wider, than the gradients from the latter two time periods. This effect appeared in all sets of data from all birds, with both single and dual S+ curves.

Working with discrete trials, Boneau, Holland, and Baker (1965) found that when responses to trials following reinforced trials were collected according to stimulus, broader gradients resulted than on other trials. Comparable gradients were constructed by locating all stimulus presentations after the appearance of a reinforced S+ and adding up (for each stimulus) the responses emitted during such trials. Since random stimulus sequences were used, the resulting sums were corrected by dividing each by the number of times the given stimulus actually appeared after a reinforced trial. The same procedure was followed for stimuli separated from reinforced presentations by two intervening trials. The resulting gradients did not differ in any systematic way from each other or from the overall gradients, though small differences may have been obscured by the relatively high variability of the curves. It is possible that this result differs from that of Boneau *et al.* (1965) because the extended trials allowed reinforcement effects to dissipate within the reinforced

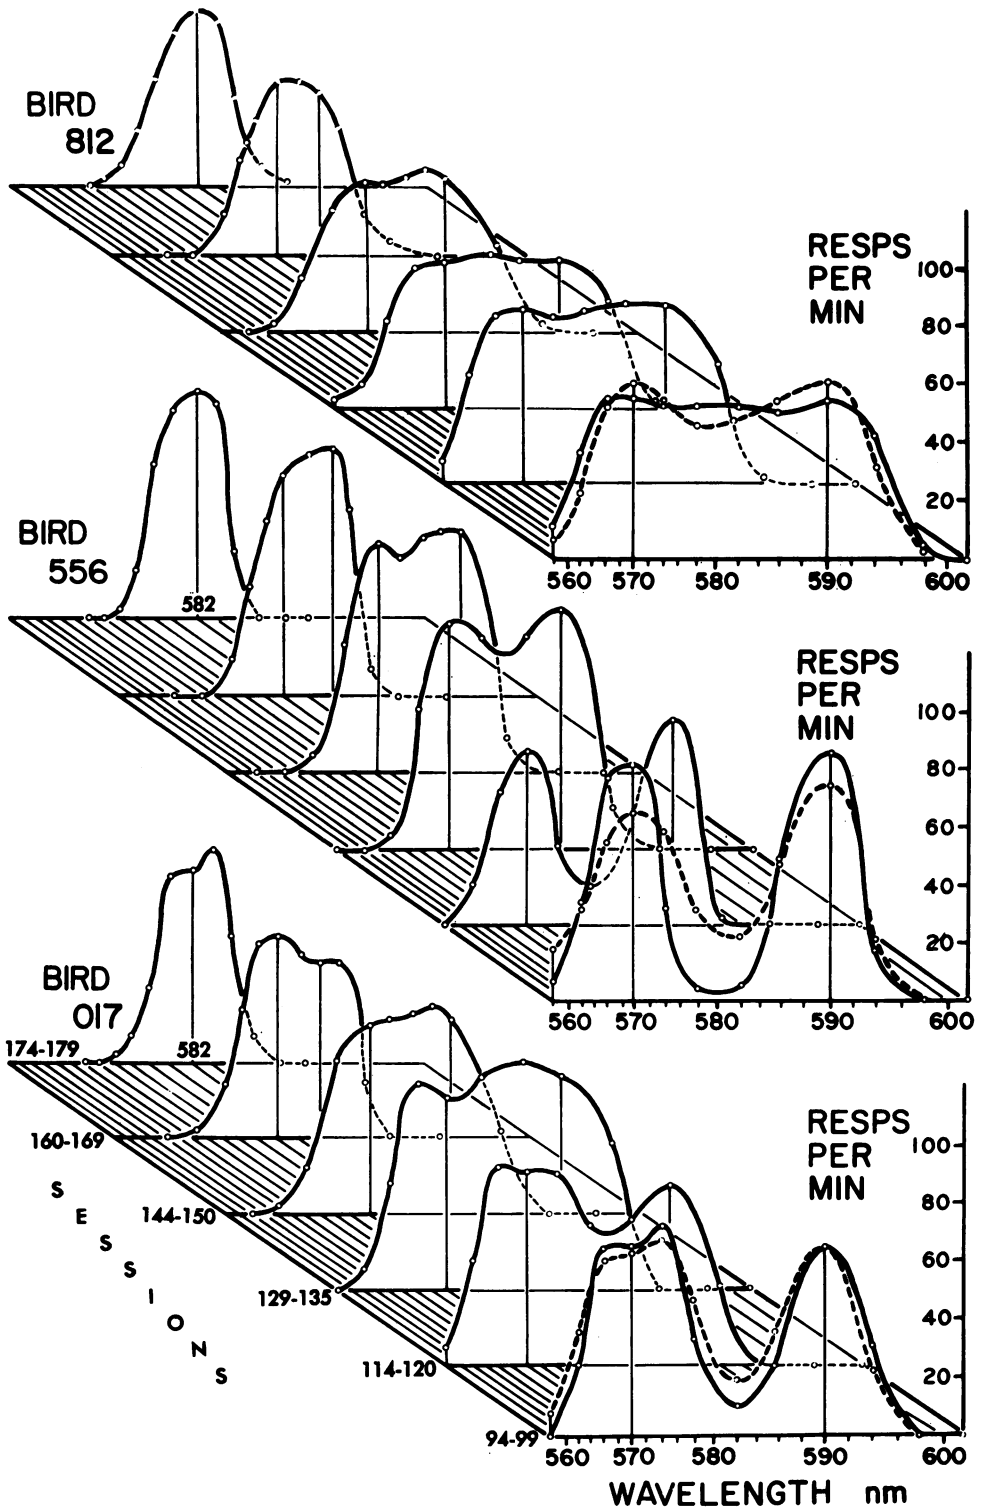


Fig. 2. Summation gradients for various S+ separations. Thin vertical lines indicate reinforced stimuli; dashed curves are replications. Note that the abscissa is distorted; see text for explanation.




Fig. 3. Sample data showing effect upon gradient shape of time within each 30-sec stimulus presentation. Note that gradients from the first part of the presentation are broader than those from later parts. These curves are representative; all data sets showed the effect.

trial itself, rather than extending across trials. Boneau *et al.* used 2-sec trials and the effect had disappeared by the second trial after reinforcement.

Finally, the probability of response as a joint function of stimulus wavelength and of time since preceding response was estimated. On the basis of previous results (Blough, 1963 and unpublished data) it is to be expected that responses terminating IRTs of less than 1 sec will be less affected by stimulus value than by longer IRT responses. The procedure of inserting a 0.6-sec blackout after each peck eliminated many of these short IRTs. The remaining short IRTs were compared with longer IRTs in the following way. The IRTs/Op statistic, which estimates the probability of response conditional upon IRT (Anger, 1956), was computed for responses to each stimulus. For this purpose, IRTs were divided into four bins: 0.6 to 1 sec, 1 to 2 sec, 2 to 4 sec, and greater than 4 sec. The results for all sets of data were similar, but the largest number of meaningful points came from the sessions involving 1-nm stimulus spacing. These were

single-stimulus gradients around 582 nm, and they appear in Fig. 4. Where few or no responses occur to a stimulus, IRTs/Op becomes quite variable, and such points are omitted. Also omitted are points for the last IRT bin (IRTs greater than 4 sec), since this transformation always yields "1" in the last bin. It is evident here that the stimulus controlled response probability effectively only in the 2- to 4-sec range, and even there the curves are quite flat compared with overall response rate curves. (The overall response rate curves appear in Fig. 1, center column, black points.)

Why are the overall rate curves steeper than the conditional probability functions in Fig. 4? The effect is largely a result of the fact that at stimuli relatively distant from S+, a rather high proportion of trials yield very few or no responses. Since at least two responses are required to define an interresponse time, trials yielding 0 or 1 response contribute nothing to the curves in Fig. 4, while trials with few responses contribute relatively little. On the other side of the coin, trials to these stimuli that do happen to yield more responses, and hence contribute to Fig. 4, are a portion of the population of trials selected on the basis of high momentary response probability. Blough (1967) found this variability in rate to different trials with the same stimulus useful in analyzing stimulus control in these situations; no more will be said of the matter here.

The data of Bird 017 deserve special mention, since they differ in several respects from those of the other birds. Particularly in evidence are irregular and flat-topped gradients (Fig. 1 and 2). The difficulty in this case can probably be traced to the nature of the bird's behavior and its interaction with the LF schedule of reinforcement. This bird tended to start each trial with a burst of responses (Fig. 3). Therefore, to meet the LF schedule criterion of balanced IRT distribution, the bird was required to "wait" to an excessive degree in the remainder of the trial. This constituted, in effect, a differential-reinforcement-of-low-rate (DRL) component in the reinforcement schedule. Such a contingency is known to flatten generalization gradients (Hearst, Koresko, and Poppen, 1964) and also steady-state gradients (Gray, 1966). It is possible that this flattening is due to stimulus control of "waiting" as a response separable from peck-

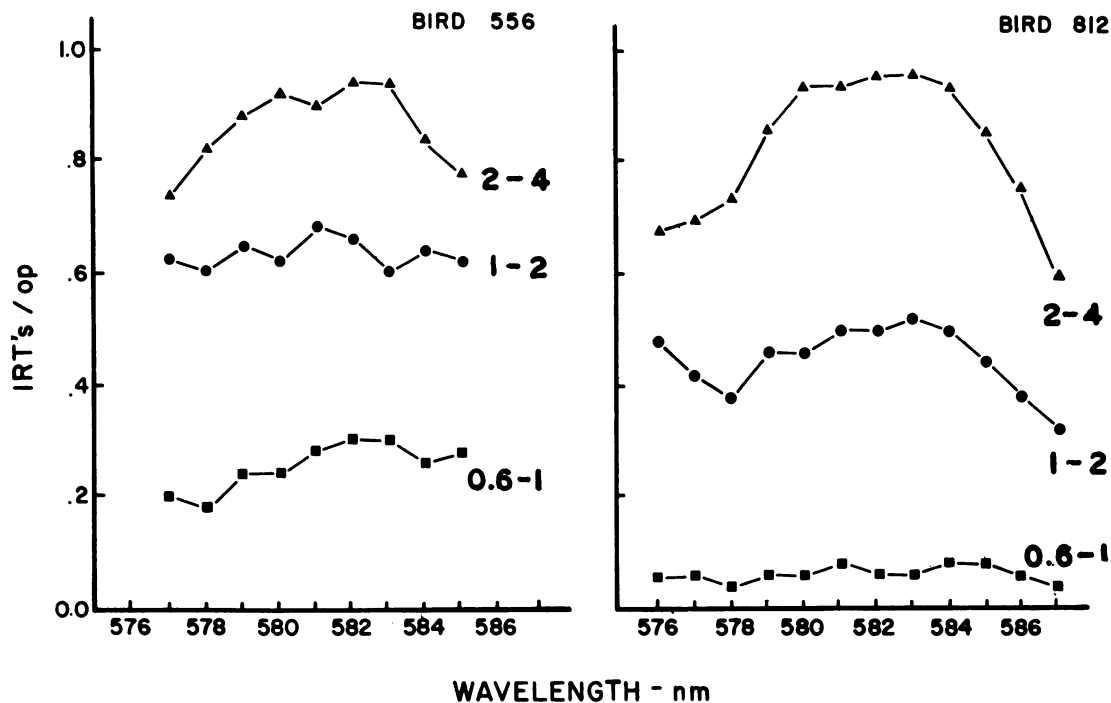


Fig. 4. Gradients showing probability of response as a joint function of wavelength and of interresponse time. The numbers next to each curve indicate the IRT bins in seconds. The overall level of the curves varies with the size of the time bin contributing to the curve; the shape of the curves is their significant aspect. See text.

ing, though further evidence is needed on this point. In any case, it is evident that the LF schedule needs modification to take into account response sequences and stimulus periodicity, as well as overall IRT distribution.

If different response patterns were found to be correlated with each of two reinforced stimuli, "summation" would be hard to interpret. Figure 2 indicates that where reinforcement was delivered in the presence of two stimuli, the rates to these two stimuli tended to remain about the same. Response patterning to the two stimuli as indicated by distributions of IRTs likewise showed no indication that different patterns of response were differentiated.

DISCUSSION

It is hard to compare these data with previously collected generalization functions, because the present data come from a situation that involved maintained reinforcement instead of extinction. As suggested above, it appears easier to cope with data of the present sort, for several reasons. First, the major variables controlling the gradient seem reasonably

well in hand, while a number of studies have suggested important effects of prior (often uncontrolled) experience upon the typical generalization gradient (*e.g.*, Friedman and Guttman, 1965; Peterson, 1962). Secondly, the maintained procedure substitutes the steady state for the complexities attendant on a transient process. In the usual generalization test, the changes due to extinction interact with those attributable to stimulus change. It is also likely that effects upon responding of stimulus novelty as such are separable from generalization effects as usually conceived. Finally, the steady state enables us to define gradients rather precisely within subjects, and to minimize the variability common to most generalization data.

Gradient shape and summation. These maintained gradients differ from typical operant generalization functions in being considerably narrower, as one would expect from the repeated presentation of reinforced stimuli against a background of unreinforced stimuli. Some readers may be struck by the rounded modes characteristic of these gradients, having been accustomed to think in

terms of the "tent-shaped" gradient popularized in textbooks. It is interesting to note that a large proportion of published generalization gradients are also best characterized by rounded modes. Where this is not the case, the apparent sharp point often arises from a lack of data points near the mode; by convention, the single high mode is then connected by straight lines to the lower points, yielding a sharp peak. The shape of the gradient at the mode is of interest with regard to summation and discrimination hypotheses. For example, if tent-shaped gradients describe both the "excitation" and "inhibition" functions, Hanson's "peak shift" (1955) cannot be derived from simple summation of the curves. Rounded modes, however, do permit prediction of the shift.

As for summation, Fig. 5 shows the relation between dual and single gradients for two birds. In this figure, the single S+ gradients (thin lines) are literally empirical curves only in the bottom row. For the other rows, which involve S+s never used singly, the curve used is the mean of the bird's three single gradients (Fig. 1) after rescaling the abscissa (as in Fig. 2). Where necessary, these single gradients were multiplied by a constant to match their peaks to the peaks of the dual S+ gradients from Fig. 2. It is, of course, impossible to collect gradients in a single S+ experiment in a manner entirely comparable to that used in a double S+ situation. In the present case, the single S+ appeared as a reinforced stimulus twice as frequently as did either of the S+s in the summation sessions. The fact that the outer limbs of the summation gradients are generally similar to the corresponding limbs of the single gradient (Fig. 5) suggests that this difference was not crucial.

All birds in all dual S+ conditions showed "summation" in the sense that the height of the dual S+ gradient exceeded that of the "component" gradients between the S+ wavelengths. There was no evidence of "summation" on the outer limbs of the gradients, a point that has been noted before (Mednick and Freedman, 1960, p. 192). There was no suggestion of any mode between the S+ wavelengths, as might be implied by theoretical accounts of a summation process (see below). In certain cases, though, the inter-S+ curve is much higher than one might expect from an additive process. For example, the individual

gradients for Bird 812 around 570 and 590 nm both fall near zero at 580 nm, yet when these wavelengths are reinforced in the same session the responding at 580 nm almost equals that at the reinforced stimuli (lower right, Fig. 5).

Theory. It is now time to examine possible combination rules that will enable prediction of dual-stimulus gradients from single-stimulus gradients. Four such rules have been stated. Three of them are mentioned by Guttman and Kalish (1956). The first states that the dual S+ gradient can be found simply by superimposing the two single S+ gradients and tracing their outline; that is, it assumes that there is no interaction or summation of the effects of dual stimulus reinforcement. Figure 5 makes it evident that this rule does not fit the present data; the portions of the dual S+ gradients between the S+s are almost all much too high.

The second rule states that the dual S+ gradient can be computed by adding the values of single gradients at corresponding points on the abscissa. This rule is equally untenable from the present data; the dual gradients are too high for widely spaced S+s and too low for closely spaced S+s.

The third rule is the "exponential addition" proposed by Hull, and derives from his Postulate 5, Corollary I: "All effective habit tendencies to a given reaction, whether positive or negative, which are active at a given time summate according to the positive growth principle exactly as would the reinforcements which would be required to produce each." (1943, p. 199) Because habit strength exponentially approaches a "physiological maximum" as asymptote, increments of strength from different sources combine to yield something less than their algebraic sum. How much less the combined strength will be depends upon the free parameter that specifies the asymptote of habit strength. Guttman and Kalish (1956) assume that response rate varies linearly with response strength, and they set the limit at 180 responses per minute. However, no matter what value is assigned to the limit, Hull's rule generates predictions at odds with the present data. A high limit is clearly untenable, since closely spaced S+ values would then generate a peak in the dual curve lying between the modes of the single gradients. A low limit could account for the

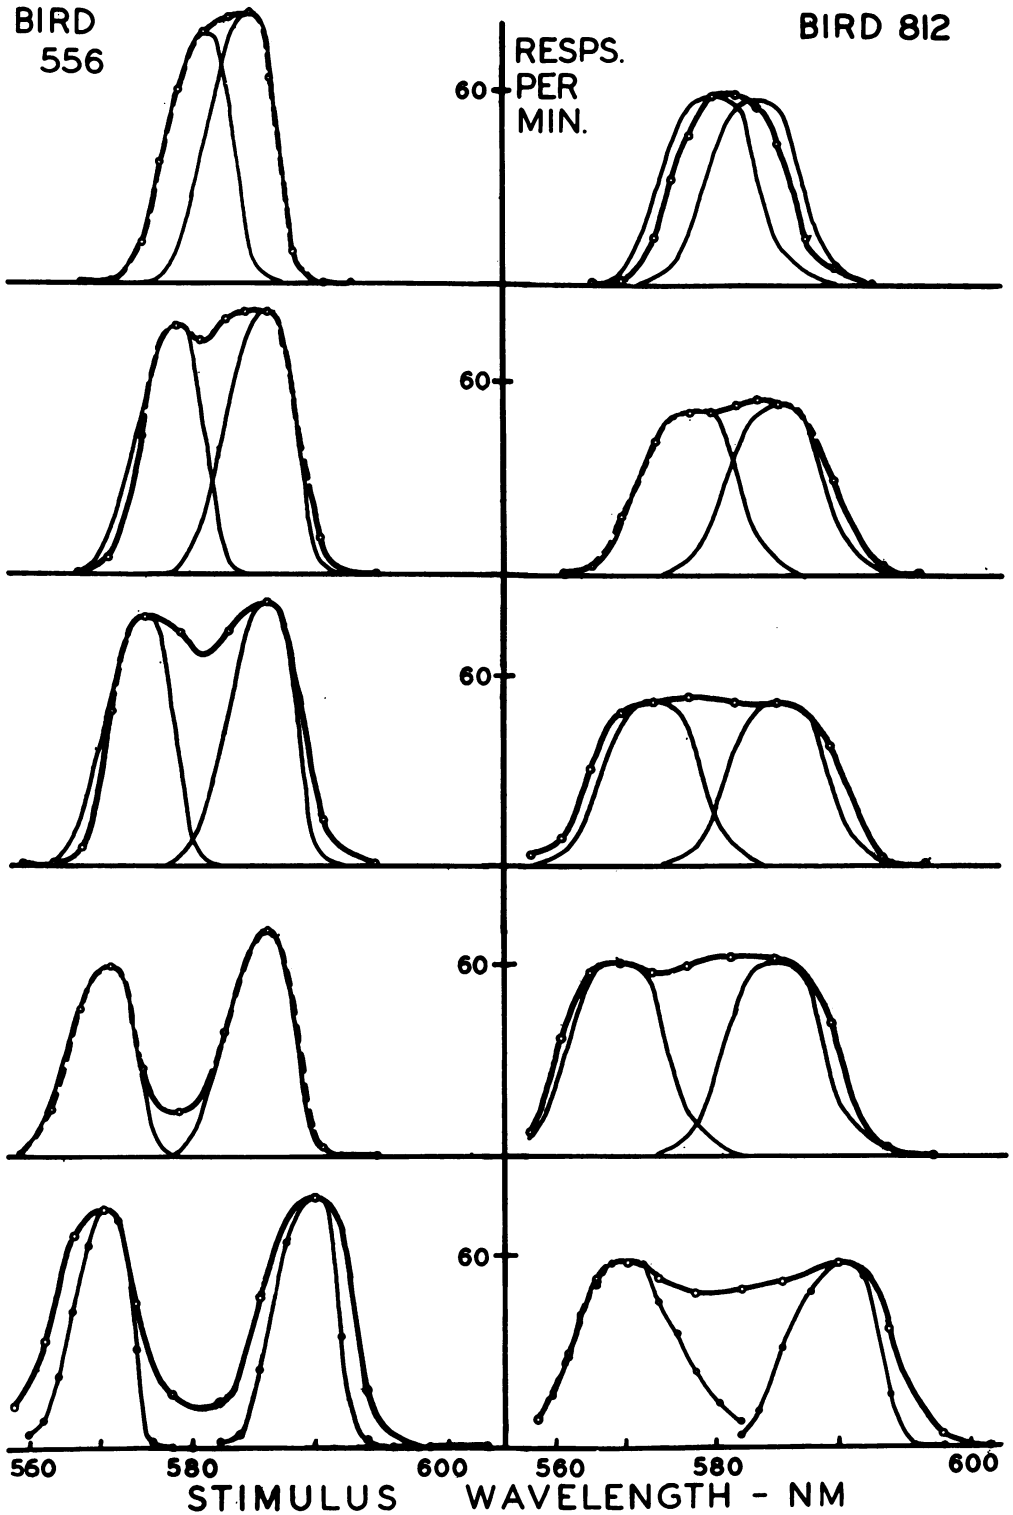


Fig. 5. Data reproduced from Fig. 2 showing the relationship of empirical "summation" gradients to single gradients derived from the same bird's data (thin lines).

lack of such peaks, but would still not account for the height of dual gradients between widely spaced S+ values (the prediction is too low) or the height of the lateral tails on dual gradients with closely spaced S+ values (the prediction is too high).

The fourth proposal for summation comes from stimulus sampling theory (Carterette, 1961; LaBerge and Martin, 1964). Here it is assumed that gradients will be linear on a properly scaled "substitutive" continuum, and summation reduces to an algebraic addition of the single S+ gradients up to a maximum given by the gradient peaks. The stimulus continua employed, the response measures used, the methods, and the subjects all differ in important ways from these items in the present research. However, insofar as the analysis can be applied to the present data, it seems inappropriate. Individual highly reliable gradients are not linear, nor will any scale transformation applicable to several gradients make them so. Also, as already mentioned, algebraic addition does not predict the present data, even when an upper limit is imposed.

It appears, then, that no formulation previously proposed adequately accounts for the present results. One could argue that no theory of "true generalization" could be expected to apply to the present data, since the steady-state procedure used does not constitute a generalization test in the usual sense of the word. This argument has a hollow ring, for there is in fact no current formulation that deals adequately with phenomena of stimulus control, call them what you will, either of the present sort, or from the more common generalization tests in extinction—such phenomena as the peak shift (Hanson, 1959) or gradient sharpening in extinction (Friedman and Guttman, 1965).

Statistical decision theory offers a possible approach to the present data and also to other findings. Though its broad implications will not be explored in detail here, it appears to offer a framework within which a variety of generalization data can be fitted with reasonable comfort. Boneau and Cole (1967) recently applied decision theory to a steady-state situation like that used in the present experiments. Let us briefly review their argument and see how it applies to generalization data.

Decision theory supposes that variations in response arise from decisions as to whether

the reinforced stimulus is present or absent. That is, when the subject sees stimulus values close to the S+, it sometimes acts as if the stimulus were the S+ and sometimes as if it were not. (In contrast, the Hullian view attributes variations in response to differences in response strength associated with each of the several test stimuli.) A physical stimulus is said to arouse within the organism a corresponding "discriminal process". Given a constant stimulus, this process varies somewhat. For differing wavelengths we might imagine, as a heuristic device, that the discriminal process corresponds to differing "hues". The probability that for a given stimulus the process takes on various values ("hues") is given by the "discriminal distribution". Figure 6 (top) suggests the discriminal distributions corresponding to two stimulating wavelengths. Criteria divide the continuum on which the discriminal process varies into regions. If, at a given moment, the state generated by a stimulus falls within the "response" region, the bird will perform that behavior which is reinforced in the presence of the discriminative stimulus. If the state falls outside the response regions, this behavior will not be emitted. In Fig. 6, 590 nm is the reinforced stimulus, and this wavelength generates states that almost

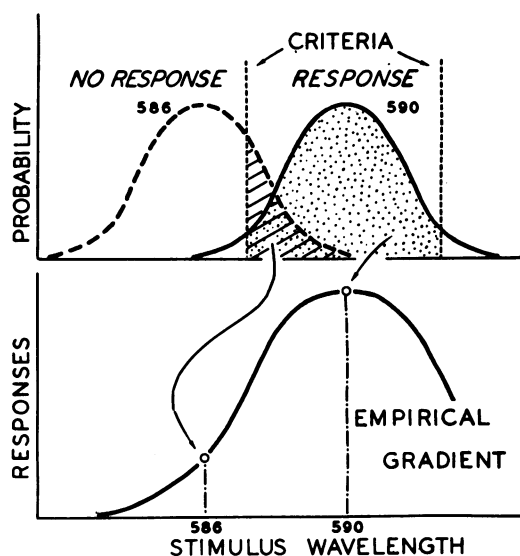


Fig. 6. Diagram suggesting how response criteria partition discriminal distributions of the several stimuli to generate the empirical gradient. Response rate to a given stimuli depends on the proportion of its distribution falling within the "response" region.

always fall within the criterion limits associated with responses. The proportion of the area under the curve that falls within these limits corresponds to the proportion of the time that this wavelength will be called "S+" and responded to accordingly. Figure 6 suggests that as the physical stimulus retreats from $S+ = 590$ nm, this proportion falls, and hence the number of responses emitted to the more distant wavelengths falls accordingly. One such case is shown, that of 586 nm. Its discriminational process falls within the criterion limits only about one-fourth of the time, and hence responding will be proportionately lower than at 590 nm.

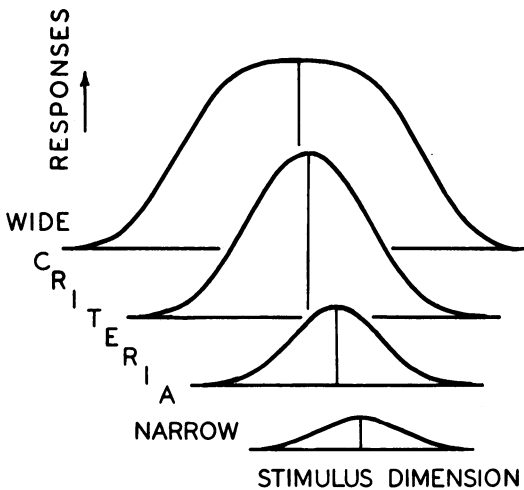


Fig. 7. Diagram suggesting how empirical gradient would be expected to vary as a result of shifts in response criteria. A Gaussian discriminational distribution is assumed here; other assumptions yield somewhat differing shapes, but all are broad and flat for wide criteria.

This way of viewing generalization gradients can handle variations in gradient width and shape as outcomes of criterion change. Figure 7 shows the gradients that result from varying the inter-criterion distance, on the assumption of Gaussian discriminational distributions. The gradients are flat for widely separated criteria, and they contract to a Gaussian form that gradually falls in height as the criteria encroach upon the discriminational distribution associated with the S+. Boneau and Cole show how conditions of reinforcement may be expected to determine an animal subject's criteria by affecting the relative payoff to be expected from various decisions. Variables such

as amount and schedule of reinforcement, deprivation, punishment, and extinction might reasonably be expected to operate in this way. Shifted peaks and asymmetries may arise if, under some conditions, the criteria on the two sides of S+ shift independently.

Detectability analysis handles summation as a simple extension of the single gradient situation. We assume two "response" regions, one associated with each of the reinforced stimuli. A stimulus between these two S+s may generate a discriminational process that sometimes falls between the two regions, and hence generates no response, and sometimes falls within one of the regions, and hence generates responses. As the S+s come closer together, intermediate stimuli will more often fall within the response region of one or the other S+. At some point, the zone between the response regions vanishes. At this point, the empirical gradient loses its twin peaks and flattens on top.

The present data seem reasonably consistent with this formulation, if one assumes rather widely spaced criteria. The summated gradients are no higher than single gradients, and when their bimodality vanishes they become flat instead of having a mode between the S+s (Fig. 2). If payoff conditions were such that response regions were very narrow, and the discriminational distribution of the S+ itself extended much beyond the criteria on either side, the theory predicts a peak between the S+s in the summation case. It would be interesting to test this prediction by manipulating payoff in the summation situation. We have already noted that summation between widely separated modes is greater than a simple algebraic sum of the single participating gradients (Fig. 2 and 5). This is also consistent with the theory, because, as the S+s approach one another, their associated response regions not only approach but also widen to meet each other. This happens because payoff determines the placement of response criteria, and the probability of reinforcement of "hues" between the S+s goes up somewhat as the S+s approach. Details of the relevant calculations can be found in Boneau and Cole (1967). Unfortunately, the expected enhancement of summation is not large enough to account entirely for the data.

Conclusion. The major contribution of decision theory to psychophysics is often said to lie in its ability to separate effects of motiva-

tional or payoff variables from effects of stimulus or sensitivity variables. The present data do not permit a "test" of the decision model, in that a separation of these variables is not possible. With no independent estimates of distribution shape or criterion locus, it is possible to say only that the data seem more compatible with the model than with the alternatives now available. But this application of decision theory represents current changes in thinking about stimulus control that are considerably more far-reaching than such a specific test. The classic view of stimulus control, exemplified in the thinking of Pavlov, held implicitly that each event distinguished by an experimenter is a "stimulus". If such a stimulus evokes a response, yet is not itself an eliciting stimulus, and has never been associated with reinforcement, its evocative power must have "spread" or generalized somehow from stimuli that were reinforced. The current view holds, of course, that environmental events are classified by organisms; if they evoke responses, it is because they fall into the class of "reinforced stimuli". The focus thus shifts from association and the spread of associative connections to the selection and classification of stimuli.

Given this reorientation, it seems worthwhile to apply the ideas of decision, criterion, and payoff to behavioral transients such as the common generalization testing procedure. For example, the narrowing of gradients during the test (Friedman and Guttman, 1965) would be expected, since a decrease in payoff should narrow the inter-criterion range. If the theory seems to make sense, we can expect experiments directed to finding what variables affect criteria in stimulus control situations. We may also expect modifications in the theory. For example, we will probably want to attribute a large measure of variance not to discriminial distributions but rather to moment-to-moment changes in criteria. It can be shown (Swets, 1964, p. 396) that models assuming criterion variance are analytically equivalent to models assuming discriminial distribution variance. Allowing criterion variance would, however, suggest that we seek variables controlling criterion variance as well as criterion locus. With this additional degree of freedom, such things as variations in summation, very broad gradients, and asymmetrical gradients could be accommodated.

Even if such speculations prove fruitful, it is clear that decision theory will not provide a complete picture of stimulus control. Discriminations are still learned somehow; the response tendencies involved may compete, be subject to inhibition, and so on. Many situations involve multiple responses that are chained or interact in complex ways; responses signifying detection may vary in rate or topography; observing behavior may be highly significant. Just as "learning theory" has dealt inadequately with stimulus problems, psychophysics has little to say of response and associative processes.

REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, 52, 145-161.
- Blough, D. S. Spectral sensitivity in the pigeon. *Journal of the Optical Society of America*, 1957, 47, 827-833.
- 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. The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 581-591.
- Blough, D. S. Stimulus generalization as signal detection in pigeons. *Science*, 1967, 158, 940-941.
- Boneau, C. A. and Cole, J. L. Decision theory, the pigeon, and the psychophysical function. *Psychological Review*, 1967, 74, 123-135.
- Boneau, C. A., Holland, M. K., and Baker, W. M. Color discrimination performance of pigeons: effects of reward. *Science*, 1965, 149, 1113-1114.
- Carterette, T. S. An application of stimulus sampling theory to summated generalization. *Journal of Experimental Psychology*, 1961, 62, 448-455.
- Clark, W. A. and Molnar, C. E. The LINC: a description of the Laboratory Instrument Computer. *Annals of the New York Academy of Sciences*, 1964, 115, 653-668.
- Friedman, H. and Guttman, N. A. A further analysis of the effects of discrimination upon stimulus generalization. In D. I. Mostofsky (Ed.), *Stimulus generalization*. Stanford: Stanford Univer. Press, 1965. Pp. 255-267.
- Gray, V. A. *The stimulus control of temporally spaced responses*. Master's thesis, Brown University, 1966.
- Guttman, N. and Kalish, H. I. Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 1956, 51, 79-88.
- Hanson, H. M. Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology*, 1959, 58, 321-334.
- Hearst, E., Koresko, M. B., and Poppen, R. Stimulus generalization and the response-reinforcement con-

- tingency. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 369-380.
- Hull, C. L. *Principles of behavior*. New York: Appleton-Century, 1943.
- Kalish, H. I. and Guttman, N. Stimulus generalization after equal training on two stimuli. *Journal of Experimental Psychology*, 1957, 53, 139-144.
- Kalish, H. I. and Guttman, N. Stimulus generalization after training on three stimuli: a test of the summation hypothesis. *Journal of Experimental Psychology*, 1959, 57, 268-272.
- LaBerge, D. and Martin, D. R. An analysis of summated generalization. *Journal of Experimental Psychology*, 1964, 68, 71-79.
- Mednick, S. A. and Freedman, J. L. Stimulus generalization. *Psychological Bulletin*, 1960, 57, 170-200.
- Morse, W. H. and Skinner, B. F. A second type of superstition in the pigeon. *American Journal of Psychology*, 1957, 70, 308-311.
- Peterson, N. Effect of monochromatic rearing on the control of responding by wavelength. *Science*, 1962, 136, 774-775.
- Pierrel, R. A generalization gradient for auditory intensity in the rat. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 303-313.
- Shepard, R. N. Approximation to uniform gradients of generalization by monotone transformations of scale. In D. I. Mostofsky (Ed.), *Stimulus generalization*. Stanford: Stanford Univer. Press, 1965. Pp. 94-110.
- Sidman, M. *Tactics of scientific research*. New York: Basic Books, 1960.
- Swets, J. A. (Ed.) *Signal detection and recognition by human observers*. New York: Wiley, 1964.
- Switalski, R. W., Lyons, J., and Thomas, D. R. Effects of interdimensional training on stimulus generalization. *Journal of Experimental Psychology*, 1966, 72, 661-666.

Received 19 January 1968.