

Tutorial

Stimulus Control: Part I

James A. Dinsmoor
Indiana University

In his effort to distinguish operant from respondent conditioning, Skinner stressed the lack of an eliciting stimulus and rejected the prevailing stereotype of Pavlovian “stimulus–response” psychology. But control by antecedent stimuli, whether classified as conditional or discriminative, is ubiquitous in the natural setting. With both respondent and operant behavior, symmetrical gradients of generalization along unrelated dimensions may be obtained following differential reinforcement in the presence and the absence of the stimulus. The slopes of these gradients serve as measures of stimulus control, and they can be steepened without applying differential reinforcement to any two points along the test dimension. Increases and decreases in stimulus control occur under the same conditions as those leading to increases and decreases in observing responses, indicating that it is the increasing frequency and duration of observation (and perhaps also of attention) that produces the separation in performances during discrimination learning.

Key words: stimulus, discrimination, generalization, gradients, observing

In the analysis of behavior, a great deal of emphasis has been placed on the control of responding by stimuli that follow the response (e.g., reinforcing stimuli), but comparatively little attention has been given to control by stimuli that precede the response (e.g., discriminative stimuli). Control by stimuli that are already present before the response occurs plays a much more significant role in everyday life than it is currently accorded in our research or even in our textbooks.

It may be that the neglect of this topic can be traced back to Skinner’s efforts to distinguish between the type of conditioning that he had begun to study in the laboratory and the older type that had previously been studied by Pavlov (Skinner, 1937, 1938). When a dog was to be trained to salivate to a conditional stimulus, as in Pavlov’s laboratory (Pavlov, 1927/1960), the response was already under a form of stimulus control. In a hungry dog, salivation could regularly and reliably be induced by presenting some dried biscuit or meat powder or a mild

solution of acid. No prior training was necessary. Because few special conditions or restrictions seemed to be attached to this method of eliciting salivation, Pavlov referred to the stimulus as an unconditional stimulus, the response to that stimulus as an unconditional response, and the relation between the two as an unconditional reflex.

When some other stimulus, like a light or a tone, was repeatedly presented shortly before the meat powder, the dog began to salivate to the new stimulus prior to the delivery of the food. As the effectiveness of this stimulus was conditional upon the pairing procedure, Pavlov called the new stimulus a conditional stimulus, the new response a conditional response, and the new reflex a conditional reflex. (In English translations, the words “conditioned” and “unconditioned” have often been used, but the original meaning is clear; see Pavlov, 1927/1960, p. 5.) Furthermore, because in this type of conditioning the food was paired with a stimulus, Skinner gave it the label *Type S conditioning*.

When it was bar pressing that was to be conditioned, however, there was no unconditional stimulus available to

Send correspondence to James Dinsmoor, Department of Psychology, Indiana University, Bloomington, Indiana 47405.

elicit the response. Insertion of the lever into the animal's chamber was a necessary but not a sufficient condition for pressing to occur. The experimenter had to wait for the animal to perform the desired action before the pellet of food could be delivered as a reinforcer. Because in this case the food was paired with a response, Skinner called it *Type R conditioning*.

Skinner also distinguished two types of behavior that corresponded to the two types of conditioning. Obviously, Pavlov's procedure could be applied only to behavior that could be elicited, prior to the training, by a specific stimulus. Because it occurred in response to a stimulus, Skinner called this type of behavior *respondent*. But for behavior like bar pressing, no equivalent stimulus could be found. Because it operated on the surrounding environment to produce the reinforcing consequence, he called this form of behavior *operant*. At the beginning of training, then, respondent behavior was characterized by an eliciting stimulus and operant behavior by its absence. During the initial process of conditioning, the only stimulus that seemed to be of any importance for operant behavior was the one that came after the response and caused it to be repeated on future occasions (reinforcing stimulus).

Although contemporary evidence suggests that Pavlovian conditioning is broad in its application and complex in its functional relations (e.g., Rescorla, 1988), this was not commonly recognized in the early years of Skinner's work. Perhaps it was to contrast his views with the prevailing stereotype of Pavlovian conditioning and to insulate himself against the charge of adopting an overly atomistic approach that Skinner rejected the label of a *stimulus-response* psychologist. In his later writings he continued to stress the relevance of other factors to operant behavior and to downplay the importance of the stimulus (e.g., preface to the 1966 printing of *The Behavior of Organisms*).

Nonetheless, even in the 1937 paper on the two types of conditioning he found it necessary to note that "it is the nature of [operant] behavior . . ." that after the initial conditioning "discriminative stimuli are practically inevitable" (1972, p. 491). And in a natural setting stimulus control is always present. It can be seen at any time, anywhere we look. All behavior is under the exquisitely detailed control of surrounding stimuli, some impinging from outside the organism, others arising from within its boundaries.

Although it is not always obvious in the way we talk and write about the subject, behavior does not occur as random strings of unrelated responses but in organized sequences, called *chains*, in which each successive response produces the stimuli, internal or external, that determine what comes next. And when the time arrives to initiate a new sequence of behavior, that, too, is signaled by a change in stimulus conditions.

If we were to monitor the stream of behavior issuing from a single individual, moment by moment throughout the day, control by antecedent stimuli would be much easier to identify than control by the organism's history of reinforcement. One form of control lies in the present, the other in the past. Another way of describing the relation between the two forms of control is to say that stimulus control is a transmission device: It is nature's way of bringing past reinforcement to bear on current behavior. To highlight the point one more time, ask yourself how you would arrange for someone seated next to you to pass the salt. Would you reinforce successive approximations, or would you present an appropriate stimulus?

DISCRIMINATION TRAINING

Although Skinner was not the first person to study the formation of a discrimination, it was his work, more than that of any other person, that brought the process into experimental focus.

The beauty of Skinner's approach lay in its simplicity. He stripped the process of discrimination down to its bare essentials. The minimum requirements were two stimuli, so that there was some difference to discriminate, and a single response, which could vary in its probability or some other characteristic. In his initial experiment, Skinner (1933, 1938) turned on a light inside the conditioning chamber and waited for the rat to press the lever. When it did so, Skinner reinforced the response by delivering a pellet of food. Next, he turned the light off for 5 min, during which time no more presses were reinforced. The cycle was then repeated, with reinforcement of a single response each time the light was turned on and subsequent extinction in the dark. Eventually, the rat was observed to press the bar promptly whenever the light came on but not to press very often when the light was not there. To demonstrate that the relation between the stimulus and the response was a product of the experimental contingencies and not of some intrinsic connection between a particular stimulus and the pressing response, Skinner also trained some of his animals on the reverse discrimination, reinforcing the pressing that occurred in the dark but not the pressing that occurred in the light. In accord with this reversal of the relation between stimulus and reinforcer, the rats learned to press the lever when in the dark but not when in the presence of the light.

Discrimination in Applied Settings

To develop the ability to apply these concepts to new situations, it is necessary to put aside the concrete stimuli and responses used in the conditioning laboratory and to concentrate on something more abstract, the pattern of events. It is the relations among the stimuli and the responses that define concepts like discrimination training, observing, and stimulus generalization. The process of learning to identify these patterns in new settings may be

facilitated by the consideration of some illustrative examples.

In 1941, when the United States entered World War II, the Army Signal Corps was faced with the task of training a large number of new recruits as quickly as possible to receive and to send messages in Morse code. That is, the trainees had to learn to transcribe sequences of long and short auditory signals (commonly represented in print by dashes and dots) into letters of the alphabet and to produce such sequences by operating a switching device. Called upon to devise an efficient training program, Fred S. Keller, himself a former telegrapher, recognized that learning to read the incoming signals was a problem in discrimination training. Ultimately, the trainees had to respond rapidly and accurately to a number of auditory patterns that sounded very much alike to the untrained ear. He proposed a technique known as the "code-voice method" (Keller, 1943), which was based directly on the procedure Skinner had used with the albino rat.

Up to that time, telegraph and radiotelegraph operators had begun their learning of Morse code by memorizing the alphabetical equivalents of the sequences of dashes and dots published in visual form, but Keller suspected that this form of training did little to help the learner discriminate among the different patterns of sound. He began immediately with the stimuli to be discriminated.

As is common when laboratory procedures developed with other animals are adapted for our highly verbal species, Keller trained his students on a number of responses (i.e., a number of different letters of the alphabet) at the same time. After a given signal had been presented through earphones or over a loudspeaker system, the trainee entered his (in those days, the trainees were all male) best guess as to the letter or digit it represented in a row of boxes on a standardized answer sheet. Shortly afterward, the instructor presented the correct response, coded as

Able, Baker, Charlie, and so on. If the correct answer matched what the trainee had just written down, the fact that his response was correct presumably reinforced that response. Note that this was an immediate reinforcement, acting upon a specific response to a specific signal, as contrasted with the usual educational practice of posting an overall score, which does not pinpoint individual stimuli or individual responses, on the following day.

If the trainee had not been able to come up with the correct answer, he now entered it in a second row of boxes, below that reserved for his original responses. In this case, it still provided a prompt that might help him on the next presentation of the same signal. (This effect was not suggested by Skinner's work but nevertheless appears to have empirical support, at least for brief intervals between prompt and test; Buchwald, 1969.)

From learning to translate sequences of long and short sounds into letters of the alphabet, it is but a short jump conceptually to learning to translate printed letters and groups of letters into the sounds of speech. By the time we begin learning to read, most of us are facile speakers of our native tongue; that is, the responses are well established, and all that is needed is to place these responses under the control of the visual stimuli. Even in English there is a substantial correspondence between individual letters and specific phonemes, so that the learner can often sound out the complete word, letter by letter, until the resemblance is recognized between the assembled fragments and the spoken word. At this stage, the learner may experience difficulty in distinguishing between similar-appearing letters like *b* and *d* or *p* and *q*. ("Mind your *ps* and *qs*.")

Training is also given in recognizing the word as a whole. I can still remember the difficulty experienced by my elder son, who was trained exclusively by the whole word method, when the word "jump" appeared for the first time at the beginning of a sentence.

Because the word began with a prominent capital J, he identified it as the only word he had previously seen with such a configuration, the proper name "Jack."

At a higher level of complexity, responding in accord with verbal instructions is also an example of the control of behavior by discriminative stimuli. The pigeon can readily learn this skill at an elementary level if pecking is reinforced in the presence of the stimulus "peck" displayed on the key but is not reinforced in the presence of the stimulus "don't peck." In applied situations, a series of stimuli may be employed to guide a sequence of responses. Or verbs and grammatical objects, for example, may be presented in novel combinations, leading to the production of novel patterns of behavior. This transfer of actions to different objects and similar features of human verbal behavior may underlie the intuitive distinction the beginning student is apt to draw between "understanding" the instructions in the human case and "simple conditioning" in the case of the pigeon. But fundamentally, both illustrate the same process.

Intermittent Reinforcement in Discrimination Training

Skinner's technique for studying the formation of a discrimination provided a useful model for things that happen in a natural setting, but from a technical point of view it left something to be desired. When the behavior in the presence of the positive stimulus was measured in terms of the time that it took before the rat pressed the lever (latency) and the behavior in the presence of the negative stimulus was measured in terms of the slope of the cumulative record (rate of pressing), there was no way to compare the subject's reactions to the two stimuli. Without a common measure, there was no way to calculate a single index that could express the relation between the two performances as the discrimination

developed. Even a rate of responding in the positive stimulus would not be very revealing if obtained under continuous reinforcement, because it would be determined mainly by the time required by the animal to consume the successive pellets of food. To provide a rate that could be compared with the rate in the presence of the negative stimulus, it was necessary to resort to an intermittent schedule of reinforcement (Dinsmoor, 1951). Then, by dividing the positive rate by the sum of the two rates, for example, an index could be calculated that would be independent of variations in the overall level of activity (Dinsmoor, 1952). Intermittent reinforcement also produced a response that persisted for a long time without further reinforcement, making it useful for comparisons of the subject's response to different values along some stimulus dimension, as in studies of generalization (Guttman & Kalish, 1956). From the time of these experiments, most studies of stimulus control have used alternating periods of variable-interval reinforcement in the presence of one stimulus and nonreinforcement in the presence of the other.

Respondent Conditioning as a Form of Discrimination Training

There is a close and interesting parallel between the procedure Pavlov used to train his dogs to salivate to a conditional stimulus and the procedure Skinner used to train his rats to form a discrimination. First, consider Skinner's procedure. A light (or some other stimulus) is presented. Skinner called this a *discriminative stimulus*. The rat responds by pressing the lever. Next, a pellet of food is delivered as a reinforcer. Then a much longer period follows, in which the positive stimulus is absent and there are no more deliveries of food. Skinner treated this period as one of negative stimulation and the difference in the rat's behavior as an indication of its discrimination between the two stimuli.

Now consider the procedure used in Pavlov's laboratory. A light (or other stimulus) is presented. Pavlov called this a *conditional stimulus* (CS), but it bears the same relation to the overall pattern of events as Skinner's discriminative stimulus. The dog responds by salivating. Next, food is delivered as an *unconditional stimulus* (the reinforcing agent in this type of conditioning). Then a much longer period follows, in which the conditional stimulus is absent and no more food is delivered. Pavlov treated this period simply as an interval of time between successive presentations of the conditional stimulus, but clearly it is the same thing that Skinner treated as a period of negative stimulation. Given the basic difference that distinguishes the two types of conditioning—pairing the food with a response or pairing it with a stimulus—the two procedures are as much alike as they can possibly be. The similarity is also preserved in more recent work: The critical difference between auto-shaping (a very popular Pavlovian procedure) and operant discrimination training based on an intermittent schedule is that in the operant case pecking is required in the presence of the positive stimulus to produce a delivery of the grain, whereas in auto-shaping the grain is delivered at the end of the stimulus period, regardless of whether any pecking has occurred. The conclusion seems inescapable that conventional Pavlovian conditioning procedures are necessarily and inherently to be categorized as forms of discrimination training.

It is not easy to illustrate the differential performance that develops under standard Pavlovian conditioning procedures, because experimenters working with those procedures do not ordinarily record the behavior of their subjects in the absence of the conditional stimulus. Like Pavlov, they view that part of the experimental routine simply as an interval between trials and take it for granted that nothing of any consequence is happening there. But Pavlov himself noted that

when conditioned reflexes are being established in dogs for the first time, it is found that the whole experimental environment, beginning with the introduction of the dog to the experimental room, acquires at first conditioned properties. . . . Later on, when the special reflex to a single definite and constant stimulus has appeared, all the other elements of the environment gradually lose their special conditioned significance. (Pavlov, 1927/1960, p. 115)

Moreover, some quantitative data showing the formation of a discrimination between the presence and the absence of the CS have been made available in a study conducted by Millenson and Dent (1971). With the technique known as *conditioned suppression*, a baseline rate of pressing a lever or pecking a key is maintained by a variable-interval schedule of nutritive reinforcement. The unconditional stimulus is a brief pulse of electric shock applied at the end of a period of exposure to a light or a tone, for example, and the reaction that is monitored (conditional response) is a reduction in the rate of pressing the lever or pecking the key in the presence of the light or the tone (CS). To measure this reduction, the experimenter compares the rate in the presence of the conditional stimulus with the rate during some control period, which is necessarily a period when the conditional stimulus is not present. Usually it is only some quantitative relation between the two rates (suppression ratio) that is reported as an index to the progress of the conditioning. In the Millenson and Dent study, however, the authors tracked the absolute rates for both stimuli, independently, for 80 sessions of training. Their Figure 1 (p. 129; reproduced here) shows that when the electric shock was first applied, both rates immediately dropped to much lower levels, but as training continued, the rate in the absence of the CS recovered, until it was substantially high-

er than the rate in the presence of the CS.


The divergence between the two rates plotted in this study provides an excellent illustration of the development of discriminative control during Pavlovian conditioning. Furthermore, the mere fact that the rate of responding in the presence of the CS has been reduced, in relation to the control rate, indicates that a similar discrimination has occurred in other studies of conditioned suppression. The suppression ratio used in these studies to monitor the level of conditioning is in effect an index to the progress of the discrimination, and the conditional stimulus is a discriminative stimulus. Notwithstanding Skinner's references to "eliciting" in the one case and "setting the occasion" in the other, the only distinction seems to lie in the type of behavior they control. (For a related discussion, see Hearst, 1975.)

STIMULUS GENERALIZATION

One of the most important of the observations coming out of Pavlov's laboratory was the discovery that when a tactile stimulus had been used as the CS, application of that same stimulus to other places on the dog's anatomy also elicited a certain amount of salivation. In addition, Pavlov noted that the nearer this stimulation was to the site of the original CS, the greater was the amount of salivation that was obtained. Or, as another example, "If a tone of 1000 d.v. [Hz] is established as a conditioned stimulus, many other tones spontaneously acquire similar properties, such properties diminishing proportionally to the intervals of these tones from the one of 1000 d.v." (Pavlov, 1927/1960, p. 113). This decline in the magnitude of the response with increases in the physical difference

→

Figure 1. For each of 3 rats, daily rates of pressing in the presence and the absence of a conditional stimulus ending with electric shock. (Reproduced from Millenson & Dent, 1971, copyright Erlbaum.)



from the original stimulus—or increase in magnitude with decreases in the physical difference—is what is meant by the gradient of generalization. The greater the similarity (closer along the physical dimension), the greater the response; the less the similarity (farther along the physical dimension), the smaller the response. Indeed, other things being equal, the degree of generalization between two stimuli may serve as a quantitative index to their similarity.

Gradients of generalization may also be obtained following operant discrimination training. Using a design that began, like Pavlov's, with a discrimination between the presence and the absence of a tone, Jenkins and Harrison (1960) reinforced key pecking on a variable-interval schedule during exposures to a frequency of 1000 Hz; between exposures, they extinguished pecking. Later, they conducted test sessions in which tones of seven different frequencies were presented without further reinforcement. The seven tones were presented a number of times, each time in a different order, so that no one tone was favored over any other by being tested early in extinction, while the response was still strong, rather than late in extinction, when it was relatively weak. All tones were tested equally often and for equal periods of time.

The result (Figure 2) was that Jenkins and Harrison's birds pecked most often to the original training tone of 1000 Hz, next most often to 670 and 1500 Hz, the tones closest on either side to the training tone, relatively little to 450 and 2250 Hz, and still less to 300 and 3500 Hz, the tones farthest from the training stimulus. (They pecked the least frequently when no tone at all was presented, because that was the negative stimulus in the discrimination training.) It is obvious that the rate of pecking by these birds was a function of the frequency of the tone. Colloquially speaking, the frequency made a difference. In this sense, the birds' behavior might be said to have




Figure 2. Gradients of generalization along the frequency dimension for each of 5 pigeons, following reinforcement in the presence of a tone of 1000 Hz and nonreinforcement during periods of silence. (Reproduced from Jenkins & Harrison, 1960, copyright American Psychological Association.)

come under the control of that feature of the stimulus.

Gradients of Generalization as a Measure of Control

Pavlov thought of stimulus generalization as an active process that extended the tendency to salivate from the original conditional stimulus to other stimuli, never before encountered by the animal, in proportion to their similarity to that stimulus. He even proposed a physiological basis for the behavior he observed: The excitation of the nerve cells, he argued, must travel from its point of origin across the cerebral cortex in much the same manner that ripples travel across the surface of a pond when a stone is dropped into it. Viewed from this perspective, stimulus generalization is an underlying process that generates a great variety of other phenomena. It has been used, for example, to account for transference in psychoanalytic therapy, visceral or muscular reactions to scenes from television or the movies (e.g., of someone teetering on the ledge of a tall building), addressing a daughter by the name of a younger sister, the origins of various phobic reactions, and so on. It may be preferable, however,

to think of these phenomena simply as illustrations of incomplete discrimination, rather than as the products of some internal process. A favorite illustration is the printed stimulus *recieve*, which in many readers will evoke the subvocal response "receive."

Clearly, generalization and discrimination are opposite numbers. As one increases, the other decreases, and vice versa. Although the standard textbook account follows Pavlov's conception of generalization, presenting it as the product of a hypothetical process of excitation or inhibition spreading from one stimulus to another, there is no reason to think of the two of them as contending forces battling for control of the individual's behavior. It makes greater sense to think of them as alternative descriptions—to think of the amount of generalization, or more specifically the slope of the gradient, as an index to the level of stimulus control that has been attained at a given point in the subject's training. To illustrate this point, consider what would happen if a pigeon were surgically deafened. Because it would have no way of distinguishing among tones of different frequency, it would peck at the same rate in the presence of all of them. The gradient of generalization would be as broad as possible (i.e., horizontal in slope). Similarly, if several different visual stimuli were presented outside of the conditioning chamber, so that the bird could not see them, it would peck at the same rate, regardless of which one was present at a given moment. The maximum amount of generalization would be recorded, but there is no reason to conjure up any kind of internal process to account for it. A high level of generalization merely indicates a low level of discrimination.

*Generalization Decreases,
Rather than Increases,
with Training*

If food is delivered in the presence of a stimulus but not in its absence, the

gradient of generalization about that stimulus becomes narrower, rather than broader, as the training continues. For example, Hearst and Koresko (1968) plotted average gradients for different groups of pigeons after 2, 4, 7, or 14 days of training. During the training, a white line was projected from the rear on an otherwise darkened key for periods of 30 s, interspersed with periods of 10 s when nothing was projected on the key. That is, the entire surface of the key was dark. During those periods when the line was displayed, pecking was reinforced on a variable-interval schedule, but when the key was completely dark, no grain was delivered. Then the birds were tested with lines of eight different tilts, ranging from the vertical to the horizontal. After only 2 days of training, a modest amount of pecking occurred in the presence of each of the test stimuli, with relatively little difference among them (Figure 3). The gradient was extremely broad. After longer periods of training, however, the number of responses as a function of position along the tilt dimension began to bunch up at the value used as the positive stimulus and to drop off more and more steeply toward the sides. Generalization was decreasing and stimulus control was increasing as the training progressed.

Although there are occasional findings in the literature on stimulus control that raise questions of interpretation, the important factor determining whether gradients of generalization become steeper with continued training appears to be the presence or absence of a correlation in time between the presence of the stimulus and the delivery of the reinforcer. Differential reinforcement with respect to a stimulus leads to steeper gradients on either side of that stimulus. Nondifferential reinforcement leads to a flattening and broadening of those gradients. There are a number of experiments that support this rule (see Dinsmoor, 1985), but the most dramatic is the one conducted by Lyons and Thomas (1967). First the authors gave their birds a series of ses-




Figure 3. Mean gradients of generalization for tilt of line displayed on key following 2, 4, 7, or 14 days of training with food deliveries in the presence of a vertical line but no food in its absence. (Reproduced from Hearst & Koresko, 1968, copyright American Psychological Association.)

sions in which pecking was reinforced in the presence of green illumination of the key but was not reinforced in the presence of a white line on a dark background. The gradient of generalization along the wavelength dimension became relatively steep. Then the same birds were given a series of sessions in which the reinforcer was delivered without regard to which stimulus was displayed on the key. The gradient became flatter. The procedure was alternated four more times, and on each occasion all of the birds showed steeper gradients following correlated reinforcement and shallower gradients following noncorrelated reinforcement. By the sign test, six determinations for each of 9 birds yields such a consistent result with a chance probability of one over the 54th power of two.

Observing the Stimulus

Early theories of discrimination learning (e.g., Skinner, 1938; Spence, 1936) attributed the divergence in performance to successive increments in strength produced by reinforcement of the response in the presence of the positive stimulus and successive decrements produced by extinction in the presence of the negative stimulus. They offered no explanation, however, as to why those increments and decrements did not continue to transfer between the two stimuli, as they did at the beginning of training, and they do not explain why the gradients of generalization should become steeper in experiments like those we have just been considering. After all, although differential reinforcement was applied to the presence versus the absence of

the positive stimulus, it was not applied to any two points along the test dimension. Where is the differential strengthening between those two points? The only plausible explanation for the growing difference in response to different points along the test dimension is that under this kind of training the subject learns to observe (e.g., look at) and presumably to attend to that and similar stimuli.

To repeat a point raised earlier, note that neither the subject in an experiment nor the ordinary person going about his or her daily business can discriminate between stimuli with which he or she has no sensory contact. To take an extreme case, a person who is currently located in London cannot react to the color displayed on a traffic signal located in Tokyo. Similarly, even if the signal is within the range of his or her vision, this same person cannot react to it unless he or she looks in the right direction to see that signal. The rays of light never impinge on the cells of the retina. Accommodation of the lens may also be necessary to bring the image into a sharp focus. Reading glasses or a magnifying device may be necessary to discriminate fine detail. A deaf person turns to the speaker to read the speaker's lips. In the supermarket, a shopper cannot judge the flavor of a newly offered brand of candy or cookies unless he or she is given a sample to place in his or her mouth. Until then, the constituent chemicals never reach the taste buds. To detect the presence of a herd of wildebeest, the lion sniffs the wind, bringing the odorous molecules into contact with the sensory membranes within its nasal passages, and to detect the approach of the lion, the wildebeest do the same. A dog approaches and sniffs the body of a person or another dog. We judge the thickness and texture of a piece of cloth by rubbing it between two fingers, we detect fever in a child by holding a hand against his or her forehead, and we sniff the carton of milk in our refrigerator to determine whether it is still fresh or has become soured. As a col-

lective category, the responses that bring about these contacts between stimulus energies and receptor cells are known, after Wyckoff (1952), as *observing responses*.

In the laboratory, the natural observing responses are difficult to record. Instrumentation to determine where the pigeon is looking from moment to moment would be extremely cumbersome, and fixing the position of the bird's head would interfere with its ability to peck the key. Experimental contingencies can be set up, however, that permit other, more readily recorded responses to mediate the subject's contact with the discriminative stimuli. Because these substitute responses have the same function as natural observing responses, they are classified as artificial observing responses. For purposes of experimental control, it is desirable that such responses have no effect on the scheduling of primary reinforcement, and Wyckoff made this a part of the standard definition. In his experimental work, Wyckoff (1969) used an electric circuit controlled by a pedal lying on the floor of the pigeon's conditioning chamber. In combination with other switching circuitry, this circuit controlled the color displayed on the key. When the bird stood on the pedal (observing), the key was either red or green, depending on whether the schedule was one of intermittent reinforcement or one of no reinforcement; when the bird was off the pedal (not observing), however, the key was always white, regardless of the schedule of reinforcement. Although the color of the key was sometimes affected by the schedule of reinforcement, note that the schedule of reinforcement was never affected by whether the bird was on or off the pedal: It was controlled by an entirely separate timing circuit.

Wyckoff considered this arrangement to be equivalent to natural behavior like looking at the key and seeing the red or the green or looking elsewhere (e.g., the wall of the chamber) and not seeing a discriminative stimulus. He found that when the red pre-




Figure 4. Operations record showing nontermination of positive discriminative stimulus but prompt termination of negative stimulus by Pigeon 5609 during last half of 45th session of discrimination training. Upward displacement of the pen indicates the presence of the event in question. (Reproduced from Dinsmoor, Thiels, Lee, Pfister, & Dougan, 1989, copyright Elsevier.)

dicted intermittent reinforcement and the green predicted no reinforcement, his pigeons stood on the pedal much more of the time than when the red and

the green were placed on a schedule that was independent of the food. He concluded that stimuli that are correlated with the presence and the absence

of a schedule of primary reinforcement become conditioned reinforcers and therefore strengthen and maintain the response that produces them.

Later studies have refined this conclusion: We now realize that it is the stimulus that is positively correlated with the primary reinforcement (S+) that becomes a conditioned reinforcer and maintains the observing behavior, whereas the stimulus that is negatively correlated with the primary reinforcement (S-) actually becomes aversive to the subject. (For reviews of the observing literature up to their respective dates of publication, see Dinsmoor, 1983; Fantino, 1977.) As an illustration, consider a detailed record of stimuli and responses obtained by Dinsmoor, Thiels, Lee, Pfister, and Dougan (1989) during the second half of one of their experimental sessions (Figure 4). The pigeon could turn a discriminative stimulus on by pecking a key to the right of the food key (observing) and could turn it off by pecking one to the left (equivalent to turning away from the stimulus). The record shows that when the stimulus that came on as a result of a peck on the observing key was positive, the pigeon left it on indefinitely; but when the stimulus turned out to be negative, within the next few seconds the pigeon usually turned it off.

In most studies of the formation of a discrimination, the experimenter does not record the subject's observing behavior. But any time reinforcers are delivered in the presence of one member of a pair of stimuli and not in the presence of the other, the conditions have been met for such behavior to be acquired. Presumably this behavior is, in fact, acquired and leads to an increase in the frequency and/or duration of the subject's sensory contact with the relevant stimuli. It is the increasing contact with the stimuli that is responsible for the increasing slope of the gradient of generalization.

There have been a number of experiments in which there seems to be no other explanation (see Dinsmoor,

1985). In these experiments, pigeons have been trained to discriminate between the presence and the absence of a specific stimulus (e.g., line, pattern, wavelength) displayed on the key and have then been tested with variations in some feature of that stimulus. Although they did share one point in common—the positive stimulus itself—the dimension along which the test stimuli were arranged was otherwise completely independent of (i.e., perpendicular to) the dimension separating the stimuli used in the training. Nevertheless, the slope of the gradient became steeper with continued training. Recall, for example, the previously cited experiment by Hearst and Koresko (1968). Pecking was reinforced in the presence of a vertical line, but no food was delivered in its absence (blank key). Then lines of several different tilts were presented during the test period, without further reinforcement. With the exception of the vertical line, note that the pigeon had never seen any of these stimuli before and had not been trained to distinguish among them; that is, pecking had never been increased by reinforcement or decreased by nonreinforcement in the presence of any other member of the series. The increasing slope of the gradient could not be attributed to differential reinforcement of responding to lines of different tilt. It was based on a characteristic of the stimulus, the degree to which each tilt in the series resembled the one used in the training, but that characteristic had been brought to bear by training applied to a dimension quite independent of the one on which the test was conducted. Something had been carried over from the training on one dimension to the testing on the other, and the most plausible suggestion is that what had been transferred was the behavior of looking at the relevant stimulus.

Gradients of Generalization Around the Negative Stimulus

In most experiments of this sort, it has been the presence of the critical




Figure 5. Mean gradients of generalization for tilt of line displayed on key following 1, 2, 4, 8, or 16 days of training with food deliveries in the absence of a vertical line but no food in its presence. (Reproduced from Farthing & Hearst, 1968, copyright American Psychological Association.)

stimulus that has served as the S+ and its absence that has served as the S-. But the experimenter can just as readily reverse that assignment and use the absence of the stimulus as the S+ and its presence as the S-. Again, if (but only if; see Honig, 1966; Weisman & Palmer, 1969) the training procedure involves continued alternation between the two stimuli and their consequences, a gradient of generalization can be obtained that reflects the similarity of each of the test stimuli to the negative stimulus.

The shape of such a gradient, however, is very different than that of a gradient around the positive stimulus (see Figure 5). Although the rate of responding to each of the test stimuli still depends on its distance from the training stimulus, instead of a peak at the value presented during the training there is now a trough. The graph representing rate of responding as a function of locus along the stimulus dimension now resembles a U or a V. Because they reflect the influence of the negative stimulus, these upside-down gradients are often called *inhibitory* gradients. The negative stimulus yields

the minimum rather than the maximum rate of responding, and the subject responds at increasingly higher, rather than increasingly lower, rates as a function of the difference between the test stimulus and that point on the stimulus dimension. Like the gradients based on the positive stimulus, this type of gradient also becomes narrower and steeper with continued training (Farthing & Hearst, 1968), again reflecting the increase in control with increasing observation of or attention to the stimulus.

One of the interesting features of the gradients obtained with negative stimuli is that they are shallower from top to bottom than the gradients obtained with their positive counterparts. Several controlled comparisons are available, and in each case the same result is obtained (see Dinsmoor, 1985): The slope of the gradient that rises on either side of the negative stimulus is never as steep as that of the gradient for the positive stimulus. That, of course, is what is to be expected, given the selective nature of the observing process. As has already been noted, when the subject in an observing experiment encounters the positive stimulus, it leaves it on, but when it encounters the negative stimulus, it turns it off. As a result, the subject receives much less exposure to the training stimulus when that stimulus has been paired with non-reinforcement than when it has been paired with reinforcement.

Combined Positive and Negative Gradients

The simple, symmetrical gradients of generalization that are obtained following special training procedures where only the positive or only the negative stimulus lies on the eventual test dimension are the exception in studies of stimulus control. A more typical arrangement is one in which both the positive and the negative stimuli lie on the same continuum. The subject learns to discriminate between two different frequencies, for example,




Figure 6. Mean gradients of generalization along the wavelength dimension following discrimination training with 550 nm as the positive stimulus (E), 560 nm as the negative stimulus (I), or both stimuli (D). (Reproduced from Marsh, 1972, copyright American Psychological Association.)

two different tilts, two different intensities, or two different wavelengths. This procedure combines the type of training that leads to a peaked gradient about the positive stimulus and the type of training that leads to a U-shaped or a V-shaped gradient about the negative stimulus. The gradient that results from this combination of the two types of training should reflect both the distance of each of the stimuli from the S+ and its distance from the S-. But because the range of variation in the subject's behavior is greater for the positive than for the negative stimulus, S+ will have a greater influence on the combined gradient, and the peak will be more obvious than the valley.

The similarities and the differences among the three types of gradient may be seen in a figure published by Marsh (1972), who obtained all three under comparable conditions (Figure 6). The stimulus dimension is the wavelength of the illumination projected on the

key. The pigeons in Group E (for excitation) were trained to discriminate between a positive wavelength of 550 nm and a negative stimulus of a plain white key (mixture of wavelengths, absence of hue). Their gradient reached a relatively sharp peak at 550 nm and trailed off on either side with increasing distance from that stimulus. The pigeons in Group I (for inhibition) were trained to discriminate between the white key as the positive stimulus and 560 nm as the negative stimulus. Their gradient was shallower but reached a minimum at 560 nm and rose on either side of that stimulus. Group D (discrimination) received the benefit of training with both 550 nm as the positive stimulus and 560 nm as the negative stimulus. In this group, the rate of pecking to a given test stimulus was affected both by the difference of that stimulus from 550 nm and by its difference from 560 nm. In effect, this

gradient combined the effects of the other two types of training.

The quantitative relations among these three types of gradient have been examined most carefully in a study conducted by Hearst (1968, 1969). For one group of pigeons, Hearst used a vertical line as his positive stimulus to obtain gradients of excitation along the tilt dimension; for a second group, he used a line of a different tilt as his negative stimulus to obtain gradients of inhibition; and for a third group, he used both of these stimuli to obtain a combined gradient. Altogether, he replicated this same basic experimental design three times, using three different degrees of separation between S- and S+. There were nine groups in all. After correcting for the possible effects of behavioral contrast, Hearst concluded that there was a relatively satisfactory fit between the sum of the gradients obtained with separate S+ training and S- training and the gradient obtained following combined training with both of these stimuli. This finding implies that it is appropriate to account for the development of the combined gradient, following normal discrimination training, with the same principles as those required to explain the changes in the positive and the negative gradients that made it up. That is, discrimination between two stimuli lying along the same continuum reflects the combined influences of the subject's observing of the S- and its observing of the S+.

IMMEDIATE DISCRIMINATION

Early theorists like Skinner (1933, 1938) and Spence (1936) saw such conventional measures of "strength" as latency, rate, proportion of responses, and resistance to extinction separate during discrimination training. When, as was often the case, there was a progressive increase in S+ and a progressive decrease in S-, it was natural for them to conclude that reinforcement of the response in the one stimulus and its extinction in the other were the factors responsible for the separa-

tion of the two performances. Yet under other circumstances, where there is a high level of observing from the beginning of the discrimination training, the separation is already evident as soon as the response is shaped, before any extinction has occurred. No strengthening or weakening seems to be required.

Consider the case in which the two stimuli are alternative colors displayed on the pigeon's key. As the color is dispersed evenly over the entire area of the key, occupying a large visual angle from the bird's point of view, it would seem to be difficult for the subject to miss seeing it as it pecks that key. In initial training conducted in my laboratory prior to a discrimination based on punishment (Dinsmoor, Flint, Smith, & Viemeister, 1969), we found that after pecking had been shaped and intermittently reinforced with the key lighted red, it was already very difficult to get the bird to peck that same key lighted with green. The bird stopped as soon as the green appeared and did not resume pecking during long periods of exposure to that color interspersed with further exposures to reinforcement in red. Terrace (1963) has documented a similar phenomenon, although it was based on a more complex sequence of procedures (see Dinsmoor, 1985, for further discussion). Despite Keller and Schoenfeld's dictum (1950, p. 119), extinction is not the hallmark of discrimination. The alternative schedules of reinforcement determine the nature of the final performance under either stimulus, as in what are known as multiple schedules, but the independence of the schedule effects requires a different explanation.

SUMMARY

This tutorial is addressed to readers who are not as conversant as they would like to be with basic principles in the control of behavior by antecedent stimuli. Its purpose is to provide a brief overview and a theoretical integration of some carefully selected re-

search findings. The present installment deals primarily with the acquisition of control by discriminative stimuli; a second installment will extend the coverage to more specialized topics that are also important to a systematic analysis of behavior.

The reader is warned, however, that the present analysis departs from standard textbook accounts in several respects:

1. In contrast to Skinner's rejection of "stimulus-response psychology," it is here suggested that antecedent stimuli exert a pervasive influence on operant, as well as respondent, behavior. They serve as a crucial link between current behavior and past reinforcement. They are the direct and immediate determinants of what the individual does at any given moment.

2. As Skinner (1938, pp. 176-177, 241) but few others have recognized, the procedure by which Pavlov trained his dogs to salivate to a conditional stimulus (respondent conditioning) is inherently a form of discrimination training. In this sense, all conditional stimuli are discriminative stimuli, and the distinction in terminology appears to rest more on the type of behavior that is under consideration and how that behavior is measured than on any difference in the role of the stimulus.

3. The Pavlovian view that stimulus generalization represents an active process in which new and different stimuli become effective because of the spread of excitation from one point to another is rejected. The gradient of generalization becomes narrower, rather than broader, as training continues. In the present account, discrimination and generalization are treated simply as alternative ways of looking at the same phenomenon, and the slope of the gradient is used as a way of measuring it.

4. In discrimination training, the positive stimulus serves as a conditioned reinforcer of observing behavior, leading to an increase in sensory contact with that stimulus; increased contact with the negative stimulus is an incidental but inevitable consequence

of the increase in observing. The decreasing transfer of the effects of reinforcement to behavior in the presence of S- and of extinction to behavior in the presence of S+ is attributed to more frequent and more prolonged observation of the two stimuli. This refers both to the time at which reinforcement is or is not delivered and to the time at which the performance is subsequently measured. Although removed from direct experimental observation, attention may play a similar role (Sutherland & Mackintosh, 1971). The respective schedules of reinforcement interact with the increasing separation of the two performances to determine their eventual levels and temporal patterns.

REFERENCES

- Buchwald, A. M. (1969). Effects of "right" and "wrong" on subsequent behavior: A new interpretation. *Psychological Review*, *76*, 132-143.
- Dinsmoor, J. A. (1951). The effect of periodic reinforcement of bar-pressing in the presence of a discriminative stimulus. *Journal of Comparative and Physiological Psychology*, *44*, 354-361.
- Dinsmoor, J. A. (1952). The effect of hunger on discriminated responding. *Journal of Abnormal and Social Psychology*, *47*, 67-72.
- Dinsmoor, J. A. (1983). Observing and conditioned reinforcement. *Behavioral and Brain Sciences*, *6*, 693-704.
- Dinsmoor, J. A. (1985). The role of observing and attention in establishing stimulus control. *Journal of the Experimental Analysis of Behavior*, *43*, 365-381.
- Dinsmoor, J. A., Flint, G. A., Smith, R. F., & Viemeister, N. F. (1969). Differential reinforcing effects of stimuli associated with the presence or absence of a schedule of punishment. In D. P. Hendry (Ed.), *Conditioned reinforcement* (pp. 357-384). Homewood, IL: Dorsey.
- Dinsmoor, J. A., Thiels, E., Lee, D. M., Pfister, J., & Dougan, J. D. (1989). Selective observing: Pigeons turn discriminative stimuli on and off by pecking separate keys. In N. W. Bond & D. A. T. Siddle (Eds.), *Proceedings, XXIVth International Congress of Psychology: Vol. 6. Psychobiology: Issues and applications* (pp. 443-456). Amsterdam: Elsevier.
- Fantino, E. (1977). Conditioned reinforcement: Choice and information. In W. K. Honig & J. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 313-339). Englewood Cliffs, NJ: Prentice-Hall.

- Farthing, G. W., & Hearst, E. (1968). Generalization gradients of inhibition after different amounts of training. *Journal of the Experimental Analysis of Behavior*, *11*, 743-752.
- Guttman, N., & Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, *51*, 79-88.
- Hearst, E. (1968). Discrimination learning as the summation of excitation and inhibition. *Science*, *162*, 1303-1306.
- Hearst, E. (1969). Excitation, inhibition and discrimination learning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 1-41). Halifax, Nova Scotia: Dalhousie University Press.
- Hearst, E. (1975). The classical-instrumental distinction: Reflexes, voluntary behavior, and categories of associative learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 2. Conditioning and behavior theory* (pp. 181-223). Hillsdale, NJ: Erlbaum.
- Hearst, E., & Koresko, M. B. (1968). Stimulus generalization and amount of prior training on variable-interval reinforcement. *Journal of Comparative and Physiological Psychology*, *66*, 133-138.
- Honig, W. K. (1966). The role of discrimination training in the generalization of punishment. *Journal of the Experimental Analysis of Behavior*, *9*, 377-384.
- Jenkins, H. M., & Harrison, R. H. (1960). Effects of discrimination training on auditory generalization. *Journal of Experimental Psychology*, *59*, 246-253.
- Keller, F. S. (1943). Studies in international Morse code: I. A new method of teaching code reception. *Journal of Applied Psychology*, *27*, 407-415.
- Keller, F. S., & Schoenfeld, W. N. (1950). *Principles of psychology*. New York: Appleton-Century-Crofts.
- Lyons, J., & Thomas, D. R. (1967). Effects of interdimensional training on stimulus generalization: II. Within-subject design. *Journal of Experimental Psychology*, *75*, 572-574.
- Marsh, G. (1972). Prediction of the peak shift in pigeons from gradients of excitation and inhibition. *Journal of Comparative and Physiological Psychology*, *81*, 262-266.
- Millenson, J. R., & Dent, J. G. (1971). Habituation of conditioned suppression. *Quarterly Journal of Experimental Psychology*, *23*, 126-134.
- Pavlov, I. P. (1960). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex* (G. V. Anrep, Trans.). New York: Dover. (Original work published 1927)
- Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*, 151-160.
- Skinner, B. F. (1933). The rate of establishment of a discrimination. *Journal of General Psychology*, *9*, 302-350.
- Skinner, B. F. (1937). Two types of conditioned reflex: A reply to Konorski and Miller. *Journal of General Psychology*, *16*, 272-279.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York: Appleton-Century-Crofts.
- Skinner, B. F. (1972). *Cumulative record: A selection of papers* (3rd ed.). New York: Appleton-Century-Crofts.
- Spence, K. W. (1936). The nature of discrimination learning in animals. *Psychological Review*, *43*, 427-449.
- Sutherland, N. S., & Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- Terrace, H. S. (1963). Discrimination learning with and without "errors." *Journal of the Experimental Analysis of Behavior*, *6*, 1-27.
- Weisman, R. G., & Palmer, J. A. (1969). Factors influencing inhibitory stimulus control: Discrimination training and prior non-differential reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*, 229-237.
- Wyckoff, L. B. (1952). The role of observing responses in discrimination learning. Part 1. *Psychological Review*, *66*, 68-78.
- Wyckoff, L. B. (1969). The role of observing responses in discrimination learning. In D. P. Hendry (Ed.), *Conditioned reinforcement* (pp. 237-250). Homewood, IL: Dorsey.