Critique of S-R Conditioning Theory

Gordon H. Bower

Stanford University

Running Title: S-R Theory

Abstract

The controversy between "behavior therapy" and "cognitive behavior modification" seems a recapitulation of the controversy between stimulus-response conditioning theory and cognitive psychology. The author, a cognitive psychologist, proposes a specific version of S-R conditioning theory, and summarizes a number of theoretical and empirical objections to this theory. The basic objection is that many human performances require more abstract theoretical mechanisms to account for them than what is possible within the strictures of S-R theory. Also S-R theory provides no adequate representation of human knowledge, of the competence brought about by learning experiences. S-R theory can appear plausible here only by encowing "covert responses" and "covert stimuli" with all the properties of "idea" or "conceptual structure". Further, the "laws of conditioning" as classically formulated are shown to fail in revealing ways. Classical and operant conditioning contingencies would appear to be arrangements which provide the person with information about environmental regularities, about the correlation between events or between actions and outcomes.

As a cognitive psychologist, I was invited to address the AABT convention in Atlanta in December, 1977 (see Bower, in press). At the convention I also attended a number of papers and symposia, several of which surprised me because they revolved around basic controversies over theoretical foundations of psychology. I had not expected so much interest and enthusiasm among clinical behavior modifiers over what were academic theoretical issues. As I listened to the several symposia on the themes of "learning theory foundations of behavior therapy", I gathered the impression that people working in the behavior-change field were not generally familiar with the theoretical arguments against a stimulus-response (S-R conditioning theory of human behavior, which a number of people in the behavior therapy field appear to believe. Therefore, I wrote this paper in the hope that it might prove useful to the field to have some of those anti-S-R arguments uncounted in one accessible place.

Before beginning, however, I should note that many people pay lip service to S-R theory (or lambast it) without any definite conception of what it is. Surely, behavior modifiers should be sensitive to unwarranted associations of labels since their field as a whole has suffered some bad press and public reactions due to unwarranted inferences about "therapeutic practices" which most behavior therapists find abhorrent. For most people, "S-R theory" appears to mean little more than the belief that behavior in the primary observable datum available to the psychologist, and that behavior is often a function of the "stimulus situation" loosely-defined. But such generalities are endorsed by practically every theoretical camp in psychology, and therefore are not distinguishing features of S-R conditioning theory.

A Specific Formulation of S-R Conditioning Theory

The content of S-R conditioning theory as I understand it commits one to considerably stronger theoretical assumptions than the afore-mentioned platitudes. In order to have a definite target for our remarks, I will use a specific version of an S-R conditioning theory of behavior. I will suppose that such a theory consists of these three fundamental assumptions:

A1. The only elements required in a psychological explanation can be put into one-to-one correspondence with potentially observable elements. These elements may themselves be observable stimuli or responses, or they may be derived from such observables. Examples of derivatives would be mediating responses, covert-response-produced stimuli, and the like.

A2. The elements mentioned in A1 become connected or associated if and only if they occur contiguously in objective time or space.

A3. All observable behavior can be explained by concatenating the associative links mentioned in A2.

These three assumptions were dubbed the "Terminal Meta-Postulate" by Bever, Fodor, and Garrett (1968), and I will abbreviate it as the TMP. Various subsidiary factors such as drive-level, incentive, CS intensity, etc. are often conceptualized as affecting the "strength" of particular S-R connections (e.g., Hull, 1943). Still, the main theoretical entities that compose and make the behavior "go" are the S-R connections.

Refutation of the Terminal Meta-Postulate

Clearly, we could refute the Terminal Meta-Postulate if we could point to any example of a performance humans are capable of but which cannot be mimicked by an S-R machine restricted by the TMP. Basically, any performance will do so if, to explain it, we must postulate some internal control elements or storage elements that are neither stimuli nor responses. These are often performances controlled by lengthy "internal computations".

One such counter-example is a performance on a mirror image language, wherein a the person is taught to discriminate symbol strings whose last half is the mirror image of their first half. To illustrate, suppose the symbols c and d are used, then examples of "grammatical" strings would be cc, dccd, dcddccddcd, and so on. If S is the start symbol for a string, then the rules for generating a mirror image string are:

1) $S \rightarrow cSc$

2) $S \rightarrow dSd$

3) $S \rightarrow$ nothing

For example, to generate cddc from S, we would apply rule 1, then rule 2, then rule 3. To say that such rules characterize the grammar is not to say that people use such rules to generate or to recornize mirror-image strings. People can be trained to distinguish between instances and non-instances of mirror-image strings that are said to them. They also can be taught to generate mirror-image strings of varying lengths. The question is whether the constraints of the TMP would permit an account of such a recognition ability or string-generation ability. Let us investigate that prospect.

Suppose we let hearing c or d correspond to the two stimuli, and saying c or d correspond to the two responses. The four obvious associations are c → c, c → d, d → d, and d → c. From a unique starting symbol, S (c or d), these associations will generate all possible strings. However, it takes little thought to see that they produce far too much--ungrammatical strings as well as grammatical ones. How, then, can we restrict the generator (or the recognizer) to produce just those strings which are mirror images?

It turns out that there is simply no way to do this within the constraints of the TMP. The problem is that introducing a given symbol early in a grammatical string commits one to writing its mate at an indefinitely later point in the sequence, and this commitment and several others must be simultaneously remembered as a number of intervening symbols are generated. Such postponed commitments result from "center embedding"; the mirror-image language is center embedded in the extreme. Embedding introduces dependencies between elements at arbitrary distances, and such remote dependencies can not be represented by associations between adjacent elements.

A system that generates (or recognizes) mirror-image strings requires something like a push-down stack memory, a device that stores symbols and retrieves them according to the principle of last-in first-out. Such a memory store itself clearly violates the TMP--for example, no unique external stimulus corresponds to "third element back in time". Let's see how such a system would generate a string like dcddcd. It would begin by generating a c and store a token of c on the stack; then generate a d and store a token of d; then it would empty the three tokens from its push-down stack, generating the last three symbols in the order dcd. The entities and operations just described all violate parts of the Terminal Meta-Postulate. For example, the tokens of c and d mentioned above violate the TMP. These c and d tokens are not "responses" in the standard sense, since standard responses are discrete time-limited events rather than representations or internal elements that reside for indefinite periods on push-down stacks in memory.

Therefore, to account for how people generate or recognize mirror image strings, we have to postulate several structures or processes which are only abstractly related to observables. This is just what the Terminal Meta-Postulate prohibits. The mirror-image example used here has no significance except to make salient the phenomenon of center-embedding rules and the difficulties they create for a surface-oriented ("all observables") view of sequential behaviors such as language. The heart of the problem is the representation of complex remote dependencies. Such remote dependencies are not at all esoteric in actual behavior; for example, the basic rule that the subject and verb of of a sentence should agree in number (singular or plural) requires a speaker-listener to keep track of just such remote dependencies.

An Experimental Analogue of Mirror-Image Performance

For the experimentally-minded, we can approximate mirror image performance in the laboratory by having subjects repeat a sub-span memory list forwards then backwards: thus, an input such as 9327 would cause the person to respond 93277239. It is difficult to imagine how one would give a strict S-R conditioning account of such a performance or, for that matter, of any new "computational performance" we instruct the subject to carry out. The response obtained to the input 9327 (or any other four digits) can be changed instantly, of course, by instructing a suitably educated subject to compute some other function, e.g., to say the number obtained from $x_1 - 2x_2 - x_3 + 2x_4$ (answer 15). There is no doubt that we can describe the computing of such a function in terms of subvocal stimuli and responses (see e.g., Staats, 1968). However, the several parts and mechanisms used in the calculations (e.g., place-holders, partial sums, scanners) are not really stimuli or responses. Also the stimulus-response language seems not to give a very cogent account of how the instructions cause the several parts of the computation-process to be set up and executed. In contrast, the idea that instructions compile directly into a mental program for operating on numbers seems a natural and heuristic way to characterize such abilities.

Types vs. Tokens

In discussing the Terminal Meta-Postulate, I touched on the idea of "tokens" as internal synbols used to represent multiple occurrences of a given stimulus "type". Karl Lashley (1951) was one of the first to point out the need for distinct instances of the same stimulus-type or response-type if sequential behaviors were to be represented properly in memory. Consider, for example, the difficulty of encoding into associations and storing the two sequences: (1) blue-red-green, and (2) gold-red-brown. If a single internal memory element, RED, is used to represent the red in both sequences and is connected via associations to its predecessors and successors, then trouble results because such a structure will incorrectly generate blue-red-brown and gold-red-green as well as the correct sequences. To keep track of the differing predecessors and successors, a standard theoretical move is to introduce the idea of tokens or copies of a given symtol-type. Tokens can be created to stand in for the type in a particular context. These copies point to or are associated to the parent type (concept) in memory, but enable different clusters of information to be associated with different occurrences of a given symbol-type. The problem is that memory tokens are neither stimuli nor responses. They violate the Terminal Meta Postulate.

Labeling of Associative Relations

Conditioning theory is a form of classical associationism, and a long-standing complaint is that it has only one type of associative connection, the S-R "bond", whereas in fact associations occur in many different types such as antecedent-to-consequence (or vice versa), ancestor-to-descendant, predecessor-to-successor, and so on. Similarly, the verb eating is related to glutton as is act to actor, to restaurant as act to location, to steak as act to object, to fork and knife as act to instruments. The alternative suggestion has been that associations between concepts should be "labeled" or differentiated in memory in some way according to their type. This would explain how the person is able to retrieve immediately from a given stimulus another item that bears a specific relation to it. For example, a person can be set to respond to each stimulus word with a second word that is opposite in meaning (e.g., up-down, left-right, male-female, etc.). The person can react very rapidly in such tests, much quicker than one would predict if large numbers of unlabeled associates of the stimulus word were being examined serially for their appropriate fit to the task demands. Accordingly, such results argue for storage of directly labeled associations.

Objections to Contiguity Law of Association

The foregoing objections are raised against either S-R theory or against associationism. More specific objections can be raised against the standard laws of association which have been adopted in nearly all conditioning theories.

Thus, it is widely believed that a <u>sufficient</u> condition that a clearly adequate stimulus becomes associated to a clearly adequate unconditioned stimulus (US) is that the CS occur contiguously with the US. But the results on <u>blocking</u> and <u>overshadowing</u> of conditioning have demonstrated that temporal contiguity of events is not sufficient. A neutral cue A may continue to remain ineffective if, when it occurs with the US, it does so in the presence of a cue B which is an already learned cue which has high predictive validity for the US (see Rescorla, 1972). Thus, for example, if cue B already predicts electric shock, then new trials of pairing the compound A + B with that shock will produce no conditioning of the A-to-shock association. Rather it appears that for a cue to acquire associative value it must add novel predictive value to whatever other cues are available in the situation.

As another example of the failure of temporal contiguity, Edward Thorndike (1931) reported several examples of invariable temporal contiguity of two events which never became associated in the person's mind because the events were not perceived as "belonging" together. To cite a simple example, Thorndike would read his subjects a very long list of full names like "Mary Jones, Bill Smith, Sam Peck, Richard Jones, Bill Smith...", repeating a number of names frequently. He found as expected that repetition of a pair increased the likelihood that the subject could recall the last name ("Smith") when cued with the first name ("Bill"). However, Thorndike found no parallel tendency for subjects to associate a last name with the <u>following</u> first-name, e.g., what first-name consistently followed Jones? Despite very high repetition frequencies for these last name-first name sequences, subjects learned nothing about them. Why? Because our linguistic habits set us to group a first-name and last-name as "belonging together" in a way that successive last- and first- names of an arbitrary sequence do not. And we rehearse and store only events that apparently belong together. Belongingness of two sensory events of attributes has also been studied as a phenomenon of perceptual organization, and a number of perceptual conditions promoting belongingness are known (see Asch, 1969). Attributes that are unified perceptually as belonging together tend to be easily "associated" as a result.

Objective Contiguity Not Necessary for Association

The studies cited suggest that temporal contiguity by itself is not sufficient to promote association. Similarly, it can be demonstrated that temporal contiguity of two objective events is not necessary to connect the two into an association. As one example, if someone burns down my house, I might recall an enemy who earlier threatened to do just that, and thereby associate the arson with him. As a second example, Garcia and Koelling (1966) demonstrated that when a rat became sick to his stomach, he associated that sickness selectively with his memory of the most recent novel food he had eaten, disregarding the myriad auditory and visual stimuli to which he had been exposed during the interval between eating and becoming sick. The rat is prepared biologically to associate a stomach upset to something it ate rather than to lights and noises in its environment. Jacoby (1974) arranged analogous laboratory conditions for promoting "mental contiguity" of events that were far apart to an input stream to a human learner. His subjects were asked to decide for each input word whether it repeated the semantic category (e.g., birds, trees) of any earlier word in the series.

His later cued recall results suggest that it is the "mental contiguity" of two events (or their memory representations), not their physical contiguity in the objective sequence, that primarily controls the extent to which the items become associated and recalled together.

Problems With the Law of Reinforcement

The law of reinforcement (or operant conditioning) asserts that when a response is followed by a reinforcer, that response will increase in its frequency. Although the principle is widely used, there are several theoretical complexities with it. First, there are problems with arriving at a valid, noncircular definition of a "reinforcer", particularly if we adopt the perspective of Premack (1965) and talk about the relative value of various activities. In Premack's terms, if A, B, C are activities thusly ordered in preference, then an opportunity to engage in A will reinforce (increase the rate of) activity B or C; similarly, B will be able to reinforce C; but B will not reinforce A. Therefore, B is not a "transituational" reinforcer. Worse, the preference ordering of the activities may be reversed for another individual or for this individual at another time (say, as he becomes satiated or deprived of different activities), with the result that the relative reinforcement of two activities may switch (see Premack, 1962). Thus, there is no non-circular definition of a reinforcer that holds up across situations. This is not a trivial issue since one does not want the law of reinforcement to be an empty tautology (i.e., things which strengthen some responses will sometimes strengthen other responses), nor do we want the law to be disconfirmed because we have mis-defined our reinforcers. One independent assessment of relative reinforcement value is simply to ask the person whether he wants A more than B; if so, his doing B should be reinforcible by restricting A's availability so that it is contingent on B. However, this independent assessment procedure requires verbally competent humans, and the presumption that people are fully cognizant of the relative reinforcement value of any two activities--and that is not always a reasonable assumption.

A related problem in defining reinforcers is that the human subject can be instructed to assign almost any relative values to different events accordingly as the social context or the experimenter requires (see Dulany, 1968). Thus, what would normally be considered to be an unpleasant event (e.g., a hot air blast in a hot room) can be converted into a desired signal for correct response by the simplest of instructions ("hot air means you're correct"). The most general statement possible is that, if at a given moment a given individual prefers event or activity A to B, then we can increase his frequency of B by withholding A from him and making it contingent on B's occurring. The only issue then is to assess preference independently.

A second stricture on the law of reinforcement is that events that normally act as reinforcers will have relatively little impact if they are presented in a context where they are not perceived as belonging to or being caused by the person's response. For example, if following a response in a verbal learning task the experimenter were to hand the subject $3.00 with the remakr, "Here are the wages I'd promised to pay you", the event would not be grouped as belonging to the specific response just made, and so that response would not be unusually strengthened. Similarly, if, as the subject responds, the experimenter's assistant asks the experimenter "You wanted sugar in your coffee, didn't you?" and the experimenter says "Right!", the eavesdropping subject does not attribute the "Right" as causally belonging to his own response. Such examples seem trivial mainly because they unveil an assumption deeply ingrained in standard contingency arrangements, namely, the social context in which the reinforcer is delivered is always one in which the reinforcer is seen as causally belonging to the operant response. The point is that reinforcers do not automatically strengthen responses they follow.

Thirdly, in human learning, presentation of the reinforcer acts as an informative event, signalling the person that he has made the correct response. Whether or not the reinforcer has some "satisfying/gratifying" effect, that affect seems to be somewhat irrelevant to the learning that the signal occasions. The clearest separation of informative versus satisfying effects of typical reinforcers (e.g., contingent money, points, praise) has occurred in experiments by Buchward (1967) and Estes (1969) using a Thorndikean trial-and-error task. Here the person chooses one of a small number of response alternatives (e.g., digits) to assign to each word of a long list of words on Trial 1. The experimenter gives certain outcomes, (Right, Wrong, or Nothing) contingent on particular stimulus-response combinations, and a test trial for learning then follows on the whole set of associations. In such an arrangement it is found routinely that S-R connections followed immediately on the initial trial by the outcome "Right" will be repeated more frequently on the test trial than are connections followed by "Wrong" or a neutral outcome. At issue is how to interpret such results. Buchwald and Estes propose a particular theoretical analysis of the situation. Suppose, to be concrete, on Trial-1 to the stimulus word SKY the subject guesses response "3" and the outcome is "Right". We suppose that at that moment he tries to establish a triplet-memory trace, "SKY-3-Right", on the basis of the contiguity and belongingness of these events. On the following trial when cued with SKY the subject tries to retrieve his previous response: if he can recall it and if he recalls the outcome "Right", he will repeat his prior response. But if he forgets the outcome, he may guess another alternative. On this account, the subject's memory of the feedback event (Right, Wrong, Nothing) regulates his tendency to repeat or shift away from the response he remembers having made to the stimulus. If this is so, then we should be able to enhance the subject's adaptive use of the reinforcing signal by delaying it long after the response on Trial 1 and not giving it in fact until the corresponding stimulus is shown on the test trial (Trial 2). Thus, when the experimenter presents SKY within the word list on Trial-2, he says to the subject "The response you made to SKY on Trial-1 was Right. How do you respond to it now?" Such delayed feedback clearly relieves the subject of having to remember feedback from the earlier trial; he need only remember his response, which event has greater likelihood than his remembering his response and its outcome. Buchward (1967) did this experiment and obtained the counterintuitive results as predicted: subjects not told Right or Wrong as they responded on Trial 1 but rather informed of this as each stimulus was presented on Trial 2 (the Delayed condition above) showed greater response-shift sensitivity to the feedback signals than did subjects given immediate feedback in the standard manner as they responded during Trial 1. That is, delayed-feedback subjects repeated Right responses and shifted away from Wrong response more than did controls. This is exactly the reverse of the outcome predicted by a "delay of reward (satisfaction)" theory. Buchwald's results and a number reviewed by Estes (1969) and Brewer (1974) demonstrate in controlled laboratory studies that human learning is largely determined by the information value of the "reinforcing signal".

Of course, there is more to reinforcement than just information--there is affect and incentive. Incentive conditions are especially relevant to influencing performance as contrasted to learning. Thus, following learning of a stimulus-response-outcome triplet, the person's selection of a given response to a stimulus will be decided according to the outcome he expects for it. Such outcome expectancies guide selection among alternative responses, and energize performance of the selected one. This is the "incentive motivation" idea that's been so prominent in Tolman's and Hull's accounts of motivation (e.g., Logan, 1960).

The Functional Response Unit

As a graduate student at Yale, I was taught my learning theory by Neal Miller and Frank Logan. I recall free-wheeling critique of S-R conditioning theory by Neal Miller in which he offered the opinion that S-R theory knew relatively little about stimuli or about responses, but was all about the "hyphen"--the connection--between stimuli and responses. He was characterizing the fact that learning theory speculated at length about the conditions necessary and/or sufficient for learning or connecting up of familiar stimuli with familiar responses, but that it provided no detailed analysis of pattern recognition or perceptual learning, and gave few guidelines for defining and isolating responses and their inter-relations (e.g., compatibilities of different responses). Frank Logan (1960) and Charles Shimp (1975, 1976) have given the most penetrating discussions of what constitutes a "response class" in learning experiments. These discussions raise many questions about the proper level at which to describe theoretically, say, a pigeon's key pecking or some such simple action. There are many variations in even simple response classes, in the way a pigeon pecks a key. Typically we aggregate these into a single response class because we (or the apparatus) do not differentially reinforce these versions. But it is always plausible to postulate some dimension of reinforcement which varies with some quantitative dimension of the response. Thus, at a minimum we can say that, on most schedules, reinforcement will be delivered sooner the more rapidly the pigeon responds. Similarly, in running to escape an electric shock, the faster an animal runs to the safety box the sooner he terminates the shock. This observation implies that different speeds of a response (or inter-response times of a pigeon's pecking) are differentially reinforced, and therefore we ought to distinguish quantitative variations as "different response classes". Furthermore, the quantitative variations in response are also differentiated by the amount of aversive work or effort they require, with faster responses usually producing more rapid accummulation of unpleasant fatigue. A theoretical analysis which assumes that quantiative variations constitute different response classes is called a "micromolar" theory. Logan and Shimp argue that even for simple responses like keypecking a micro-molar analysis of reinforcement schedules is always more revealing of uniformities than is a relatively more molar approach. We know that we can differentially reinforce fast or slow response rates, by scheduling reinforcement contingent on very short or very long inter-response times. In further developments along this line, Shimp (1975, 1976) has been able to reinforce and strengthen entire sequential patterns of inter-response times (e.g., short, long, long, short). He suggests that patterns of inter-response times (not single pecks) may be the "bird's eye view" of what are the functional response units in the time-honored Skinner box.

The Complexity of Sensori-Motor Skills

To carry the objection a step further, S-R theory has been virtually abandoned as unhelpful by modern workers analyzing human skilled performance (see, e.g., Schmidt, 1975). The components necessary to synthesize a given skill (e.g., returning a tennis serve) are appreciably more complicated than the simple reflex arc suggests. Recent theorizing (eg., Schmidt, 1975) proposes that such skills are represented by complex schema that coordinate several types of information, namely, (1) the person's posture and the location of the limb to-be-moved, (2) the computed location and time of arrival of the target (e.g., the on-rushing tennis ball), (3) specification of the parameters of coordinated motor-command sequences distributed over time and over the muscles of the body which will meet the target, (4) the sensori-motor feedback expected as the movement proceeds, and (5) the reinforcing outcome expected from successful execution of the response. The feedback from initial movements is compared to the target and is used in servo fashion to correct errors and readjust movement parameters so as to hit the target. This conception focuses much more on the continuous modulation of an activity in interaction with its environment, and departs markedly from the ballistic "knee-jerk" model suggested by classical S-R theory.

Actions Infused with Intentions

A long-standing philosophical literature (e.g., Goldman, 1970; Taylor, 1964) argues that we cannot properly describe or characterize what someone is doing by referring only to his muscle movements. Rather, we commonly describe the apparent goal or intention of an action (e.g., "He's trying to start a fire by sparking flints"), and individual movements acquire their significance by virtue of their role in that plan. The claim is that actions are not reducible to, or equivalent to, a set of muscle movements. As a practical matter, it is alsmot impossible to describe a response without reference to its goal or its attainments, or to the stimuli which the muscles are in contact with, or to the feedback provided by these contacts. In fact, the exact movements required to perform an instance of an action class (e.g., "eating a scrambled egg") must vary with the circumstances, being different for when I use a fork or fingers, when I am standing or lying down, for when my left or right arm is in a cast, etc. Furthermore, a given movement is classified as two different actions depending upon the inferred intention or "meaning" of the action in a given social context. The Godfather's "Kiss of Death" is not an amorous gesture, though his bondaire bussing of his wife probably is.

Earlier, I noted that individual movements acquire their interpretive significance by virtue of their role in a plan. You can appreciate this by imagining how some rule-governed behavior would be described differently by an observer who knows the rules and the "meaning" of prescribed forms and a second observer who is completely ignorant. To be concrete, imagine listening to a football game as described by one announcer who is knowledgeable about football versus a second announcer who knows no football. The latter might say "Player #18 took the football from #93, backpeddled 5 steps, then player #32 took the ball form him and ran forward at a 135° angle". The knowledgeable announcer would say something like "The quarterback faded back as though to pass, then handed off on a draw play to the fullback who went through the hole left by the suckered left-tackle". The difference in descriptions is one of the intentions that play significant roles in the game. "Backpeddling 5 steps" is not the same as "fading back to pass". The defense, of course, "reads" the apparent intentions of the quarterback and reacts accordingly, by moving back into a pass defense. The effectiveness of feints, cons, bluffs and other such put-ons depends upon a shared cultural system for anticipating what someone is likely to do. In competitive games, of course, it is useful to mislead your adversary to get a slight edge, for a quarterback to "fake" a line plunge before passing or vice versa, for a boxer to "fake" exhaustion before raining a fusillade of punches on his incautious opponent. Clearly, then, we interpret others' actions according to motives we infer. And, of course, someone can be said to be trying to do something even though the action never comes off successfully. A quarterback can be "fading back to pass" even though he soon falls down or gets sacked before he gets the pass off. This discussion of action concepts and their use in everyday life sounds noncontroversial until we realize that it flies in the face of a basic preaching of classical behaviorism, namely, that behavior should be described at the level of colorless muscle movements. Cognitive psychologists have always sided with common sense on this issue, claiming that intentions and

and expectancies must be taken into account in describing actions. The intentions of actions are also terribly important too for practically all legal systems, which try to adjudicate responsibility and retribution for someone's harmful actions depending upon how much he foresaw and intended the outcome. For instance, the decision between first-degree vs second-degree murder vs. manslaughter is largely one about intentions of the agent who killed someone.

Memory

Stimulus-response theory has, in my opinion, never been able to give a very cogent account of many simple facts of human memory, and the criticisms have been available for quite some time (for an early critique, see Russell, 1927; for a later critique, see Anderson & Bower, 1973, Chpt. 2). One problem is that one cannot imagine what are the hypothetical "responses" that are supposedly conditioned when a person learns to recognize a person, place, or thing. In an experiment, say, on face recognition, I show you a number of photographs of people's faces to learn; later I test you by showing you a series of some old and some new photographs, and you are to say "old" to those you recognize as having seen before. People can clearly do this, but what response is being learned here? Surely it is not the response of saying "old" or the reaction of "I've seen this before" to a given face, because in fact the person had just the opposite reaction when he saw the face earlier. Also, it can't be an "emotional" or "attitudinal" reaction to the face because we recognize and distinguish far more faces than we ever have emotional responses. A standard dodge here is to say that the person has a "perceptional response" to the initial presentation of the face, and that recognition comes about from the strength of this perceptual response during a second presentation of the face. In such discussions, "Perceptual response" is used with approximately the non-technical meaning of "image", "percept", or "internal representation", and the "response" vocabulary buys no conceptual power whatever. It is simply a token how to stimulus-response psychology.

Stimulus-response theory similarly misses on explaining judgments people can make about their memories. As one example, suppose we present some faces more often than others to the unsuspecting subject. We will find that he is able later to judge rather accurately the number of times each photograph had been shown to him. Yet, he is manifestly not "counting" during each presentation of a photograph, and it is not clear that he is making any "response" while he looks at a face a second or third time. The point is that the subject's frequency judgment hears no simple relation to any response he might have made as he looked at the face. As a second example, a person has knowledge of what he knows without being able to recall. In "tip of the tongue" experiences, people realize they know a name (i.e., could recognize it) which they are unable to recall at the moment.

A basic objection to the stimulus-response account of memory is that it confuses the knowledge a person has about some event with the performances (responses) he may use to indicate that knowledge. But the layman knows that he can ask many questions about memory of a given event, and obtain different answers depending on the question. Suppose I tell you a new fact, such as "The population of Tempe, Arizona is about 100,000 "now". As memory performances, I can ask you to rephrase that fact, to list Arizona cities with populations over 80,000, to mimic my tone of voice when you repeat that sentence, to say whether Tempe's population is 50,000, to decide whether Tempe's population is at least a quarter of Tucson's, and to decide whether Tempe qualifies for a low-income housing project from HUD. These are all performances reflecting memory of the event of my telling you the population of Tempe, yet the responses per se bear no similarities to one another. Rather they have conceptual meanings that are related to the meaning of the utterance, namely, that Tempe has a current population of about 100,000. These issues of memory representation are particularly problematical when we are discussing memory for discourse, for texts, or even for single sentences. The logical propositions (the "deep structure") underlying an utterance are often compressed and transformed in the "surface form" of the utterance; and the information in a set of propositions can be imparted and rephrased in many different ways, and it can be tapped in many di-ferent ways. Thus, for example, we recognize that two paraphrases are different renderings of the same conceptualization. Or that a sentence in Spanish means the same as one in English. In such instances, remembering a proposition is clearly not the same as remembering a verbal utterance, since one can paraphrase the content of a proposition just as one can parrot a foreign-language utterance with no knowledge of what proposition it refers to. All of this is to argue that the knowledge gained from reading or from listening to a lecture is not easily represented in terms of stimulus-response sequences.

Bandura (1971) has made a similar argument regarding the retention of a modeled action sequence an observer witnesses. There is no mechanism in simple S-R theory to encode the observed sequence, and to retain it for some time before it is imitated, perhaps in entirely different stimulus circumstances. We can consider that knowledge to be represented as associations among action concepts or images, but that is a rather different theoretical vocabulary from S-R theory.

Although I have cited the problems for S-R theory raised by event representations in memory of verbally competent humans, the problems are there for nonverbal humans and lower organisms. Thus, a nonverbal retardate can learn to imitate different complex responses at a delay, and we must describe how he can encode and retain a distinctive motor command. Similarly, monkeys, rats and pigeons can be taught to perform on delayed matching to sample, a form of short-term memory. How does the organism retain a representation of the sample stimulus over the delay interval in order to cue his later choice of the matching stimulus? An earlier hypothesis was that some immediate "adjunctive" behavior chain kicked off by the sample cue continued differentially and long enough to mediate the matching choice. However, search for the mediating chains (see, e.g., D'Amato, 1973) have proven singularly unrevealing. Mostly, animals seem to remember without "externalizing" that memory in their peripheral musculature throughout the retention interval. Rather they seem to remember the "idea" or "proposition" that the sample stimulus this trial was the green triangle. We can, of course, represent that knowledge in terms of associations between pre-existing concepts (i.e., "Current Sample" → "green triangle"), but that is a move not permitted the theorist of S-R persuasions. And that, of course, is the heart of the problem.

Final Commentary

I have tried to summarize some of the arguments against an S-R conditioning theory of human behavior. More extensive accounts of each argument are available in the primary sources cited. In addition the interested reader may profitably read a review chapter by William Brewer (1974) with the arresting title of "There is no convincing evidence for operant and classical conditioning in humans". What Brewer means is that there is no evidence that humans learn contingencies without being aware of them, that is, in a conditioning situation, they first "figure out" what is going on in the scheduling of experimental events (or something correlated to the actual events), and then they response according to what they want. Brewer argues, correctly I believe, against the hypothesis that reinforcement contingencies bring about conditioning automatically, whether or not the person is aware of them.

In my address at the AABT meeting (Bower, in press), I argued that behavior therapy should use cognitive psychology as a theoretical base because most therapeutic techniques are more readily explained by cognitive theory than by S-R conditioning theory. Certainly, such a move would prevent the frequent criticism that behavior therapy violates tenets of behaviorism (eg., Breger & McGaugh, 1965; Locke, 1971) and would blunt the stupid attacks on behavior therapy as arch "behaviorism" in the popular press.

In ending this recital of objections to S-R conditioning theory, I should point out that it is always possible to argue around one or another objection by blurring the claims of the theory, backing them into vagueness. I believe much vagueness has been introduced by moving the hypothetical stimuli and responses "inside" the organism, by such constructs as "mediating response", "covert response", "stimuli from covert responses", and so on. In such uses, all the standard properties of stimuli and responses are left behind, and the terms are used with all the surplus meaning of idea, concept, or image; consequently, the terms are used without the strictures of the Terminal Meta-Postulate. I believe (as does Fodor, 1965) that, when used in strict accord with the Terminal Meta-Postulate, mediating or covert responses provide little explanatory power over and above concepts referring to observable stimuli and responses. However, I am willing to concede that there is widespread disagreement, and perhaps confusion, on this point.

References

Anderson, J. R. & Bower, G. H. Human Associative Memory. Washington, D. C.: Winston, 1973.

Asch, S. E. Reformulation of the problem of association. American Psychologist, 1969, 24, 92-102.

Bandura, A. Psychological Modeling: Conflicting Theories. New York: Aldine-Atherton, 1971.

Bever, T. G., Fodor, J. A., & Garrett, M. A formal limitation on association. In T. R. Dixon & D. L. Horton (Eds.), Verbal Behavior and General Behavior Theory. Englewood Cliffs, N. J.: Prentice-Hall, 1968.

Bower, G. H. Contacts of cognitive psychology with social learning theory. Cognitive Therapy and Research, in press.

Breger, L., & McGough, J. L. Critique and reformulation of "learning theory" approaches to psychotherapy and neurosis. Psychological Bulletin, 1965, 63, 338-358.

Brewer, W. F. There is no convincing evidence for operant and classical conditioning in humans. In W. B. Weimer & D. S. Palermo (Eds.), Cognition and Symbolic Processes. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1974.

Buchwald, A. M. Effects of immediate vs. delayed outcomes in associative learning. Journal of Verbal Learning and Verbal Behavior, 1967, 6, 317-320.

D'amato, M. R. Delayed matching and short-term memory in monkeys. In G. H. Bower (Ed.), The Psychology of Learning and Motivation, Vol. 7, New York: Academic Press.

Dulany, D. E. Awareness, rules, and propositional control: A confrontation with S-R behavior theory. In T. R. Dixon & D. L. Horton (Eds.), Verbal Behavior and General Behavior Theory. Englewood Cliffs, N. J.: Prentice Hall, 1968.

Estes, W. K. Reinforcement in human learning. In J. Tapp (Ed.), Reinforcement and Behavior. New York: Academic Press, 1969.

Fodor, J. A. Could meaning be an r_m? Journal of Verbal Learning and Verbal Behavior, 1965, 4, 73-81.

Garcia, J., & Koelling, R. A. Relation of cue to consequence in avoidance learning. Psychonomic Science, 1966, 4, 123-124.

Goldman, A. I. A Theory of Human Action. Englewood Cliffs, N. J.: Prentice Hall, Inc. 1970.

Jacoby, L. L. The role of mental contiguity in memory: Registration and retrieval effects. Journal of Verbal Learning and Verbal Behavior, 1974, 13, 483-496.

Lashley, K. S. The problem of serial order in behavior. In L. A. Jeffress (Ed.), Cerebral Mechanisms in Behavior. New York: Wiley, 1951.

Locke, E. A. Is "behavior therapy" behavioristic? Psychological Bulletin, 1971, 76, 318-327.

Logan, F. A. Incentive. New Haven: Yale University Press, 1960.

Premack, D. Reversibility of the reinforcement relation. Science, 1962, 136 255-257.

Premack, D. Reinforcement theory. In M. R. Jones (Ed.), Nebraska Symposium on Motivation, 1965. Lincoln: University of Nebraska Press, 1965.

Rescorla, R. A. Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), The Psychology of Learning and Motivation (Vol. 6). New York: Academic Press, 1972.

Russell, B. R. Philosophy. New York: Norton, 1927.

Schmidt, R. A. A schema theory of discrete motor skill learning. Psychological Review, 1975, 82, 225-260.

Shimp. C. P. Organization in memory and behavior. Journal of the Experimental Analysis of Behavior, 1976, 26, 113-120.

Shimp, C. P. Perspectives on the behavioral unit: choice behavior in animals. In W. K. Estes (Ed.), Handbook of Learning and Cognitive Processes, Vol. 2, Hillsdale, N. J.: Lawrence Erlbaum Assoc., 1975.

Staats, A. W. Learning, Language, and Cognition. New York: Holt, Rinehart and Winston, 1968.

Taylor, C. The Explanation of Behavior. London: Routledge & Kegan, Paul, 1964.

Thorndike, E. L. Human Learning. New York: Century Co., 1931. (Paperback edition, Cambridge, Mass.: MIT Press, 1966.)

Reference Note

1. The author's work is supported by a research grant, MH-13905, from the National Institutes of Mental Health. Requests for reprints should be sent to Dr. Gordon H. Bower, Dept. of Psychology, Stanford University, Stanford, Calif. 94305.