

How to Teach Animals

B. F. Skinner

Teaching, it is often said, is an art, but we have increasing reason to hope that it may eventually become a science. We have already discovered enough about the nature of learning to devise training techniques which are much more rapid and give more reliable results than the rule-of-thumb methods of the past. Tested on animals, the new techniques have proved superior to traditional methods of professional animal trainers; they yield more remarkable results with much less effort.

It takes rather subtle laboratory conditions to test an animal's full learning capacity, but the reader may be surprised at how much he can accomplish even under informal circumstances at home. Since nearly everyone at some time or other has tried, or wished he knew how, to train a dog, a cat, or some other animal, perhaps the most useful way to explain the learning process is to describe some simple experiments which the reader can perform himself.

"Catch your rabbit" is the first item in a well-known recipe for rabbit stew. Your first move, of course, is to choose an experimental subject. Any available animal—a cat, a dog, a pigeon, a mouse, a parrot, a chicken, a pig—will do. (Children or other members of your family may also be available, but it is suggested that you save them until you have had practice with less valuable material.) Suppose you choose a dog.

The second thing you will need is something your subject wants, say food. This serves as a reward or—to use a term which is less likely to be misunderstood—a "reinforcement" for the desired behavior. Many things besides food are reinforcing—for example, simply letting the dog out for a run—but food is usually the easiest to administer in the kind of experiment to be described here. If you use food, you must of course perform the experiment when the dog is hungry, perhaps just before his dinnertime.

The reinforcement gives you a means of controlling the behavior of the animal. It rests on the simple principle that whenever something reinforces a particular form of behavior, it increases the chances that the animal will repeat that behavior. This makes it possible to shape an animal's behavior almost as a sculptor shapes a lump of clay. There is, of course, nothing new in this principle. What is new is a better understanding of the conditions under which reinforcement works best.

To be effective a reinforcement must be given almost simultaneously with the desired behavior; a delay of even one second destroys much of the effect. This means that offering food in the usual way is likely to be ineffective; it is not fast enough. The best way to reinforce the behavior with the necessary speed is to use a "conditioned" reinforcer. This is a signal which the animal has observed in association with food. The animal is always given food immediately after the signal, and the signal itself then becomes a reinforcer. The better the association between the two events, the better the result.

For a conditioned reinforcer you need a clear signal which can be given instantly and to which the subject is sure to respond. It may be a noise or a flash of light. A whistle is not recommended because of the time it takes to draw a breath before blowing it. A visual signal like a wave of the arm may not always be seen by the animal. A convenient signal is a rap on a table with a small hard object or the noise of a high-pitched device such as a "cricket."

You are now ready to start the experiment with your dog. Work in a convenient place as free as possible from distraction. Let us say that you have chosen a "cricket" as your conditioner reinforcer. To build up its reinforcing power begin by tossing a few pieces of food, one at a time and not oftener than once or twice a minute, where the dog may eat them. Use pieces so small that 30 or 40 will not appreciably reduce the animal's hunger. As soon as the dog eats pieces readily and without delay, begin to pair the cricket with the food. Sound the cricket and then toss a piece of food. Wait half a minute or so and repeat. Sound the cricket suddenly, without any preparatory

movement such as reaching for food.

At this stage your subject will probably show well-marked begging behavior. It may watch you intently, perhaps jump on you, and so on. You must break up this behavior, because it will interfere with other parts of the experiment. Never sound the cricket or give food when the dog is close to you or facing you. Wait until it turns away, then reinforce. Your conditioned reinforcer will be working properly when the dog turns immediately, and approaches the spot where it receives food. Test this several times.

Wait until the dog is in a fairly unusual position, then sound the signal. Time spent in making sure the dog immediately approaches the food will later be saved manyfold.

Now, having established the noise as a reinforcer, you may begin teaching the dog. To get the feel of the technique start with some simple task, such as getting the dog to approach the handle on a low cupboard door and touch it with its nose. At first you reinforce any activity which will be part of the final completed act of approaching and touching the handle of the cupboard. The only permissible contact between you and the dog is via the cricket and the food. Do not touch the dog, talk to it, coax it, "draw its attention," or interfere in any other way with the experiment. If your subject just sits, you may have to begin by reinforcing any movement, however slight. As soon as the dog moves, sound the cricket and give food. Remember that your reaction time is important. Try to reinforce as nearly simultaneously with the movement as possible.

After your subject has begun to move about, reinforce when it turns to the cupboard. Almost immediately you will notice a change in its behavior. It will begin to face toward the cupboard most of the time. Then begin to reinforce only when the dog moves nearer the cupboard. (If you withhold reinforcement too long at this stage, you may lose the facing response. If so, go back and pick it up.) In a very short time—perhaps a minute or two—you should have the dog standing close to the cupboard. Now begin to pay attention to its head. Reinforce any movement which brings the nose close to the handle. You will have to make special efforts now to reduce the time between the movement and the reinforcement to the very minimum. Presently the dog will touch the handle with its nose, and after reinforcement it will repeat this behavior so long as it remains hungry.

Usually it takes no more than five minutes, even for a beginner, to teach a dog this behavior. Moreover, the dog does not have to be particularly "smart" to learn it; contrary to the usual view, all normal dogs learn with about equal facility with this conditioning technique.

Before going on with other experiments test the effect of your conditioned reinforcer again two or three times. If the dog responds quickly and eats without delay you may safely continue. You should "extinguish" the response the dog has already learned, however, before teaching it another. Stop reinforcing the act of touching the cupboard handle until the dog abandons this activity.

As a second test, let us say, you want to teach the dog to lift its head in the air and turn around to the right. The general procedure is the same, but you may need some help in sharpening your observation of the behavior to be reinforced. As a guide to the height to which the dog's head is to be raised, sight some horizontal line on the wall across the room. Whenever the dog, in its random movements, lifts its head above this line, reinforce immediately. You will soon see the head rising above the line and more frequently. Now raise your sights slightly and reinforce when the head rises above the new level. By a series of gradual steps can get the dog to hold its head much higher than usual. After this you begin to emphasize any turning movement in a clockwise direction the head is high. Eventually the dog should execute a kind of dance. If you use available food carefully, a single session should suffice for setting up this behavior.

Having tested your ability to produce these simple responses, you feel confident enough to approach a more complex assignment. This suppose you try working with a pigeon. Pigeons do not tame easily. You will probably want a cage to help control the bird, and for this you can rig up a large cardboard carton with a screen or lattice top and window in the side for observing the bird. It is much less disturbing to the bird if you watch it from below its line of vision than if you peer at it from above. In general keep yourself out of the experimental situation as much as possible. You may still use a cricket as a conditioned reinforcer, and feed the bird by dropping a few grains of pigeon food into a small dish through a hole in the wall. It may take several daily feedings to get the bird readily and to respond quickly to the cricket.

Your assignment is to teach the pigeon to identify the visual patterns on playing cards. To begin with, hang a single card on a nail on the wall of the cage a few inches above the floor so that the pigeon can easily peck it. After you have trained the bird to peck the card by reinforcing the movements which lead to that end, change the card and again reinforce the peck. If you shuffle the cards and present them at random, the pigeon will learn to peck any card offered.

Now begin to teach it to discriminate among the cards. Let us say using diamonds and clubs (excluding face cards and aces) and want the bird to select diamonds. Reinforce only when the card presented is a diamond, never when it is a club. Almost immediately the bird will begin to show a preference for diamonds. You can speed up its progress toward complete rejection of clubs by discontinuing the experiment for a moment (a mild form of punishment) whenever it pecks a club. A good conditioned, punishment is simply to turn off the light ["blacking out"] or cover or remove the card. After half a minute replace the card or turn on the light and continue the experiment. Under these conditions the response is positively reinforced with food remains part of the repertoire of the bird, while the response which leads to a blackout quickly disappears.

There is an amusing variation of this experiment by which you can make it appear that a pigeon can be taught to read. You simply use two printed cards bearing the words PECK and DON'T PECK, respectively. By reinforcing responses to PECK and blacking out when the bird pecks DON'T PECK, it is quite easy to train the bird to obey the commands on the cards.

The pigeon can also be taught the somewhat more "intellectual" performance of matching a sample object. Let us say the sample to be matched is a certain card. Fasten three cards to a board, with one above and the two others side by side just below it. The board is placed so that the bird can reach all the cards through windows cut in the side of the cage. After training the bird to peck a card of any kind impartially in all three positions, present the three chosen cards. The sample to be matched, say the three of diamonds, is at the top, and below it put a three of diamonds and a three of clubs. If the bird pecks the sample three of diamonds at the top, do nothing. If it pecks the matching three of diamonds below, reinforce it; if it pecks the three of clubs, black out. After each correct response and reinforcement, switch the positions of the two lower cards. The pigeon should soon match the sample each time. Conversely, it can also be taught to select the card which does not match the sample. It is important to reinforce correct choices immediately. Your own behavior must be letter-perfect if you are to expect perfection from your subject. The task can be made easier if the pigeon is conditioned to peck the sample card before you begin to train it to match the sample.

In a more elaborate variation of this experiment we have found it possible to make a pigeon choose among four words so that it appears to "name the suit" of the sample card. You prepare four cards about the size of small calling cards, each bearing in block letters the name of a suit: SPADES, HEARTS, DIAMONDS, and CLUBS. Fasten these side by side in a row and teach the pigeon to peck them by reinforcing in the usual way. Now arrange a sample playing card just above them. Cover the name cards and reinforce the pigeon a few times for pecking the sample. Now present, say, the three of diamonds as the sample. When the pigeon pecks it, immediately uncover the name cards. If the pigeon pecks DIAMONDS, reinforce instantly. If it pecks a wrong name instead, black out for half a minute and then resume the experiment with the three of diamonds still in place and the name cards covered. After a correct choice, change the sample card to a different suit while the pigeon is eating. Always keep the names covered until the sample card has been pecked. Within a short time you should have the bird following the full sequence of pecking the sample and then the appropriate name card. As time passes the correct name will be pecked more and more frequently and, if you do not too often reinforce wrong responses or neglect to reinforce right ones, the pigeon should soon become letter-perfect.

A toy piano offers interesting possibilities for performances of a more artistic nature. Reinforce any movement of the pigeon that leads toward its pressing a key. Then, by using reinforcements and blackouts appropriately, narrow the response to a given key. Then build up a two-note sequence, reinforcing only when the sequence has been completed and by blacking out when any other combination of keys is struck. The two-note sequence will quickly emerge. Other notes may then be added. Pigeons, chickens, small dogs, and cats have been taught in this way to play tunes of four or five notes. The situation soon becomes too complicated, however, for the casual experimenter. You will find it difficult to control the tempo, and the reinforcing contingencies

become very complex. The limit of such an experiment is determined as much by the experimenter's skill as by that of the animal. In the laboratory we have been able to provide assistance to the experimenter by setting up complicated devices which always reinforce consistently and avoid exhaustion of the experimenter's patience.

The increased precision of the laboratory also makes it possible to guarantee performance up to the point of almost complete certainty. When relevant conditions have been controlled, the behavior of the organism is fully determined. Behavior may be sustained in full strength for many hours by utilizing different schedules of reinforcement. Some of these correspond to the contingencies established in industry in daily wages or in piece-work pay; others resemble the subtle but powerful contingencies of gambling devices, which are notorious for their ability to command sustained behavior.

The human baby is an excellent subject in experiments of the kind described here. You will not need to interfere with feeding schedules or create any other state of deprivation, because the human infant can be reinforced by very trivial environmental events; it does not need such reward as food. Almost any "feedback" from the environment is reinforcing if it is not too intense. A crumpled newspaper, a pan and a spoon, or any convenient noisemaker quickly generates appropriate behavior, oft" amusing in its violence. The baby's rattle is based upon this principle.

One reinforcer to which babies often respond is the flashing on and off of a table lamp. Select some arbitrary response—for example, lifting the hand. Whenever the baby lifts its hand, flash the light. In a short time a well-defined response will be generated. (Human babies are just as "smart" as dogs or pigeons in this respect.) Incidentally, the baby will enjoy the experiment.

The same principle is at work in the behavior of older children and adults. Important among human reinforcements are those aspects of the behavior of others, often very subtle, which we call "attention," "approval" and "affection." Behavior which is successful in achieving these reinforcements may come to dominate the repertoire of the individual.

All this may be easily used—and just as easily misused—in our relations with other people. To the reader who is anxious to advance to the human subject a word of caution is in order. Reinforcement is only one of the procedures through which we alter behavior. To use it, we must build up some degree of deprivation or at least permit a deprivation to prevail which it is within our power to reduce. We must embark upon a program in which we sometimes apply relevant reinforcement and sometimes withhold it. In doing this, we are quite likely to generate emotional effects. Unfortunately the science of behavior is not yet as successful in controlling emotion as it is in shaping practical behavior.

A scientific analysis can, however, bring about a better understanding of personal relations. We are almost always reinforcing the behavior of others, whether we mean to be or not. A familiar problem is that of the child who seems to take an almost pathological delight in annoying its parents. In many cases this is the result of conditioning which is very similar to the animal training we have discussed. The attention, approval, and affection which a mother gives a child are all extremely powerful reinforcements. Any behavior of the child which produces these consequences is likely to be strengthened. The mother may unwittingly promote the very behavior she does not want. For example, when she is busy she is likely not to respond to a call or request made in a quiet tone of voice. She may answer the child only when it raises its voice. The average intensity of the child's vocal behavior therefore moves up to another level—precisely as the head of the dog in our experiment was raised to a new height. Eventually the mother gets used to this level and again reinforces only louder instances. This vicious circle brings about louder and louder behavior. The child's voice may also vary in intonation, and any change in the direction of unpleasantness is more likely to get the attention of the mother and is therefore strengthened. One might even say that "annoying" behavior is just that behavior which is especially effective in arousing another person to action. The mother behaves, in fact, as if she had been given the assignment of teaching the child to be annoying! The remedy in such a case is simply well-defined response will be generated. (Human babies are just as "smart" as dogs or pigeons in this respect.) Incidentally, the baby will enjoy the experiment.

The same principle is at work in the behavior of older children and adults. Important among human reinforcements are those aspects of the behavior of others, often very subtle, which we call

"attention," "approval" and "affection." Behavior which is successful in achieving these reinforcements may come to dominate the repertoire of the individual.

All this may be easily used—and just as easily misused—in our relations with other people. To the reader who is anxious to advance to the human subject a word of caution is in order. Reinforcement is only one of the procedures through which we alter behavior. To use it, we must build up some degree of deprivation or at least permit a deprivation to prevail which it is within our power to reduce. We must embark upon a program in which we sometimes apply relevant reinforcement and sometimes withhold it. In doing this, we are quite likely to generate emotional effects. Unfortunately the science of behavior is not yet as successful in controlling emotion as it is in shaping practical behavior.

A scientific analysis can, however, bring about a better understanding of personal relations. We are almost always reinforcing the behavior of others, whether we mean to be or not. A familiar problem is that of the child who seems to take an almost pathological delight in annoying its parents. In many cases this is the result of conditioning which is very similar to the animal training we have discussed. The attention, approval, and affection which a mother gives a child are all extremely powerful reinforcements. Any behavior of the child which produces these consequences is likely to be strengthened. The mother may unwittingly promote the very behavior she does not want. For example, when she is busy she is likely not to respond to a call or request made in a quiet tone of voice. She may answer the child only when it raises its voice. The average intensity of the child's vocal behavior therefore moves up to another level—precisely as the head of the dog in our experiment was raised to a new height. Eventually the mother gets used to this level and again reinforces only louder instances. This vicious circle brings about louder and louder behavior. The child's voice may also vary in intonation, and any change in the direction of unpleasantness is more likely to get the attention of the mother and is therefore strengthened. One might even say that "annoying" behavior is just that behavior which is especially effective in arousing another person to action. The mother behaves, in fact, as if she had been given the assignment of teaching the child to be annoying! The remedy in such a case is simply for the mother to make sure that she responds with attention and affection to most if not all the responses of the child which are of acceptable intensity and tone of voice and that she never reinforces the annoying forms of behavior.