Behavior Research Methods, Instruments, & Computers
1994, 26 (2), 134-141

Sniffy, the virtual rat: Simulated operant conditioning

JEFF GRAHAM, TOM ALLOWAY, and LESTER KRAMES
Erindale College, University of Toronto, Mississauga, Ontario, Canada

We report on the use of our Sniffy program to teach operant conditioning to 900 introductory psychology students. The simulation is designed primarily to teach the principles of shaping and partial reinforcement in an operant chamber. Advanced features are provided for exploring modeling issues and the learning parameters of the model. Students observe the rat's pretraining behaviors, shape barpressing, and explore the effects of partial reinforcement schedules on a cumulative record. Any of 30 actions can be trained to occur in specific locations in the Skinner box. This paper summarizes details about the software, interface, and instructional objectives.

The Sniffy program is designed primarily to teach the principles of shaping and partial reinforcement in operant conditioning. Advanced features are provided for the exploration of modeling issues and the learning parameters of the model. This program simulates many of the behaviors one would observe in a real rat learning to operate in the controlled environment of an operant chamber. Sniffy, a simulated rat, can be trained to perform any of the 30 behaviors in its repertoire by pairing food delivery with the target behavior. So that students may fully appreciate the major features of operant conditioning, we have provided instructions for two 2-h lab sessions.

The first allows students to train Sniffy to press the bar for food, and the second explores the changes in behavior that occur under partial reinforcement (PRF) conditions. Spontaneous recovery, discrimination, and chaining are phenomema that have not yet been implemented in our program. We do not claim that the exact slope of the response rates observed in real rats have been replicated, although we have attempted to display the typical differences observed among the four major PRF schedules. Careful instructors may refer to Sniffy as *rattus siliconus* and treat the learning objective as one in which the behavior of this new species needs to documented, verified, and tested.

In the following sections, we present a rationale for the development of this simulation, a brief review of operant conditioning, and a summary of the two labs that we have run to date. Some of the background material is adapted from the program's extensive documentation. The final section provides a general description of the program and of how its major features are implemented.

We gratefully thank Erindale College, the University of Toronto Computer Shop, and Apple Canada for the support that made our teaching lab possible. Greg Wilson is the mastermind behind Sniffy's behavior, and we thank him for making the achievement of an improbable task a reality. Correspondence should be addressed to J. Graham, Psychology Department, Erindale College, University of Toronto, 3359 Mississauga Rd., Mississauga, ON, Canada L5L 1C6.

Ethics and Economics

The ideal way to learn about operant conditioning would be to work with a real rat in a real operant chamber. However, financial and ethical constraints make this impractical in most university and college settings. An operant chamber capable of delivering food reinforcement, along with a computer to control events in the chamber, record the animal's barpresses, and produce printable records would cost from $2,000 to $3,000 per student work station.

Purchasing and maintaining rats is also expensive. A young adult rat suitable for training in an operant chamber costs about $10 from a commercial distributor. In both the U.S. and Canada, animal-care regulations specify that all laboratory animals must be housed under specified conditions, that they must be kept in specially designed facilities, and that they must be cared for by specially trained animal-care technicians. Typical animal care costs $10 per rat per month. Most universities and colleges do not have facilities to house large numbers of animals used for teaching, and even if they did have enough room, the cost would be great. Thus, in recent years, few undergraduate students have been able to get hands-on experience with operant conditioning, even though operant conditioning is one of the most important topics covered in undergraduate psychology courses. Overcoming these financial barriers is one of the main reasons we created Sniffy.

Other considerations are ethical. Whether or not one is ever ethically justified to use animals in scientific research has become a hotly debated topic in recent years. Some argue that the use of laboratory animals is always unethical, but more would probably agree that the use of research animals is justified if the animals are well treated and if the research is likely to produce substantial new scientific knowledge. The use of animals for teaching purposes—where no new scientific knowledge will be gained—is harder to justify. However, experiments on operant conditioning of the type that Sniffy simulates would cause no pain whatsoever and would produce little, if any, physical discomfort to a live animal.

Simulation and Approximation

Were Sniffy a live animal, he would be a domestic laboratory rat. All domestic rats belong to the species *Rattus norvegicus*, one of the two species of rats that are common pests in buildings. Domestic rats were created in the 19th century through the selective breeding of stocks of captive wild rats. The most obvious physical difference between domestic and wild rats has to do with coloration. All commonly used domestic rat breeds have some or all of the genes for albinism, or lack of normal body pigmentation. The fully albino strains have white coats and pink eyes. The partly albino strains have patches of darker hair and normal, dark eyes.

The most obvious behavioral difference between domestic and wild rats has to do with tameness. Wild rats are normally ferocious and hard to handle. If one tried to pick up a caged wild rat, it would almost surely try to bite. Wild rats are difficult to tame even if born in captivity and handled regularly from an early age. In contrast, domestic rats are very gentle. They rarely try to bite. If treated kindly, domestic rats make interesting, affectionate, and intelligent pets.

Were Sniffy a real animal, he would have been born in captivity, in a domestic-rat breeding facility. Several large companies sell rats to laboratories, and the domestic pet trade and many universities and research institutes maintain their own rat-breeding facilities. Sniffy would probably be 90-120 days old—a young adult—at the time he was selected for training in an operant-conditioning experiment.

After selection, a real Sniffy would likely be subjected to a 2-week period of preparation for training. During this period, he would live by himself in a cage in which food and water were continuously available and every day an animal-care technician would remove him from his cage briefly and handle him gently. In this way, he would learn not to be afraid of handling. Otherwise he would be nervous and difficult to train. This standard 2-week gentling period produces fearless animals that are ready for the next phase of their education.

In operant conditioning, food is the most common positive reinforcer for training rats to press a bar. If the rat is not hungry, food is not an effective positive reinforcer. Thus, real rats are deprived of food for 24 h prior to training to make food an effective reinforcer. A major difference between Sniffy and a real rat concerns satiation. Real rats are subject to satiation for food. Sniffy is insatiable! He is always hungry, and food is always an effective reinforcer. This is one of the reasons why Sniffy is somewhat easier to train than a real rat would be.

Teaching Operant Conditioning

An animal's behavioral repertoire is often said to consist of two major kinds of behaviors: respondent behaviors and operant behaviors. Respondent behaviors are those that can be reliably elicited from untrained animals by specific, easily defined stimuli. For example, placing food in a hungry animal's mouth will elicit salivation; directing a jet of compressed air against an animal's eye will elicit an eye blink. Classical conditioning offers a set of techniques for getting new, initially ineffective stimuli to elicit respondent behaviors.

However, most behaviors do not have stimuli that will reliably elicit them. These other, "operant," behaviors are behaviors that an animal is said to emit. In an operant chamber, Sniffy or a real rat will walk around, rear up against the side walls, scratch its ears, and lick its genitals. Once in a while, it will even press the bar mounted on one of the walls. The animal does these things spontaneously. Operant conditioning is a set of related procedures that employ reinforcement (reward) and punishment to increase and decrease the frequencies of such emitted or operant behaviors.

LAB 1
Magazine Training and Shaping the Barpress

The purpose of this lab is to demonstrate basic operant-conditioning procedures, including the establishment of baseline behaviors, magazine training, shaping, acquisition (learning), and extinction (of the conditioned behavior, not the rat!). When the program is started, Sniffy is seen moving around in the operant chamber beginning at the center of the floor. In Figure 1, Sniffy is exploring the right-hand corner closest to the viewer about 8 min after being placed on an extinction schedule. On the left of the back wall is Sniffy's water spigot. Just as a real rat, Sniffy can have a drink of water whenever he wants. In the center of the back wall is the bar that Sniffy is trained to press. Directly below the bar is the hopper into which Sniffy's food pellets will drop.

At the bottom of the screen is Sniffy's cumulative record of barpresses. As time elapses, a line will be drawn horizontally across the screen from left to right. Every time Sniffy presses the bar, the line will move up a notch; every time Sniffy gets a food pellet because he pressed the bar, a little blip will be drawn across the line. When the line reaches the top of the record, it will reset itself at the bottom. If Sniffy is not barpressing, the line will be horizontal. Once barpressing has been established, the steepness of the line will reflect the rate of Sniffy's barpressing. The faster he presses, the more steeply the line will rise. The vertical lines (alternating dotted and solid) represent 5 min. In addition to these time markers, there are heavy vertical lines produced when the record resets itself.

Baseline Behaviors

Operant conditioning begins with the discovery of the untrained rat's natural behaviors. Operant conditioning affects the frequency of spontaneous behaviors, so it is important to find out what Sniffy does spontaneously. First, one simply watches what Sniffy does. Students are advised to be precise, and they are cautioned to avoid drawing inferences about what the rat likes or what he

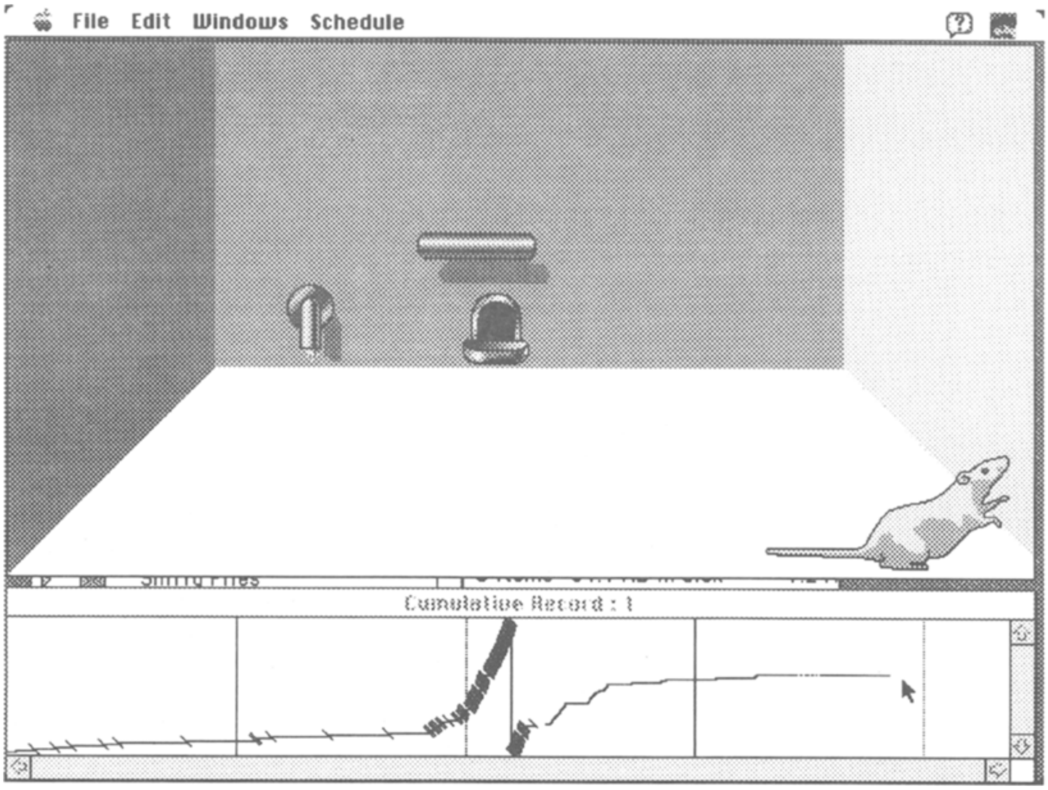

Figure 1. View of Sniffy in the operant chamber. The cumulative record shows barpressing during the acquisition and extinction of barpressing.

is thinking about. The actions must be described objectively and recorded continuously. The rat stands, licks, scratches, etc. Interval or event recording techniques can be introduced at the instructor's discretion. The primary concern is to measure the baseline rate of the response that one is attempting to condition. In this exercise, Sniffy is trained to press the bar (although it is possible to train Sniffy to perform any of the behaviors that he emits). Thus, students should record Sniffy's baseline rate of barpressing by counting the number of times he touches and presses the bar spontaneously.

Magazine Training

Sniffy is trained with positive reinforcement to press the bar more often. In this procedure, positive reinforcers are used to increase the frequency of a target behavior. Positive reinforcers are stimuli whose presentation after a target behavior makes that behavior more likely to occur again under similar circumstances in the future. In other words, positive reinforcers are stimuli that an animal will work to get. Basically, the positive reinforcement procedure consists of waiting for a target behavior to occur and then delivering the positive reinforcer. To be effective, the reinforcer must be delivered to the animal immediately after the target behavior has occurred. The immediacy with which the reinforcer is delivered is very important. If the reinforcer is delayed even a second or two, instead of reinforcing the target behavior, the reinforcer will strengthen whatever behavior the animal was performing a second or two after it performed the target behavior.

The need for immediacy of reinforcement brings to light a problem with food as a positive reinforcer. To deliver a reinforcer, students position their cursor on the bar above the hopper and click the mouse button. They observe how long it takes Sniffy to find and eat the food. Unless Sniffy is very near the food hopper when the pellet drops, he will not find the food immediately. So what does the food pellet (the positive reinforcer) reinforce? It reinforces Sniffy for the last thing he did before he ate it, which was poking his nose into the food hopper. One needs a positive reinforcer that can be delivered immediately after Sniffy presses the bar; and that is where magazine training comes in.

Reinforcers can be either primary or secondary. A primary reinforcer is a stimulus whose reinforcing power is intrinsic to the stimulus, provided that the animal is in the right physiological state. Food is a primary reinforcer for a food-deprived rat; water is a primary reinforcer for a water-deprived rat; and a sexual partner is a primary reinforcer for a rat that is ready to mate. A secondary reinforcer is a stimulus that has acquired reinforcing power as a result of being paired with a primary reinforcer. Nearly any stimulus that is not intrinsically a reinforcer can become a secondary reinforcer if it is paired with a primary reinforcer. Magazine training is the name of the procedure that one employs to turn the sound made by the food delivery mechanism into a secondary reinforcer for Sniffy.

To start magazine training, the student must wait until Sniffy is near the food hopper; then, one must operate the magazine to deliver a pellet of food. If Sniffy is close enough, he will find it quickly and start to form an association between the sound and the presence of the food pellet. To save time, several pellets of food should be given before he starts to wander off. Gradually, one begins delivering the food pellets when Sniffy is a little farther away from the hopper. He should orient to the hopper, walk over, and consume each pellet. When the student can "call" him from any part of the chamber by operating the magazine, magazine training is complete. Whenever the food-delivery mechanism is operated, Sniffy will be instantly reinforced for doing whatever he was doing just before he heard the sound.

Shaping

Sniffy has an extensive behavioral repertoire, and barpressing is part of it, even before training. Since the operant chamber is programmed to deliver a pellet of food for every barpress (the default set when the program is launched), now that he has been magazine trained, Sniffy will learn to barpress all by himself if he is left alone for an hour or so. He might even have learned it without magazine training, but it would have taken him a lot longer.

However, if one is observant and has good timing, it is possible to speed up this learning process by employing a technique called *shaping*. This procedure is employed to train an animal to do something often that it normally does rarely (or not at all), by reinforcing successive approximations of the desired behavior. Shaping an animal takes patience, careful observation, and good timing. It is a skill that one can learn with practice. Sniffy is easier to shape than a real rat would be, partly because he never gets enough to eat and partly because his behavior is not as variable as a real rat's. But he is difficult enough to shape for students to get an idea of both the frustration and the feeling of triumph that the procedure engenders.

As the first approximation of barpressing, Sniffy is reinforced for rearing up on his hind legs anywhere in the chamber. Next, once he is rearing fairly often, the student could require him to rear up against the back wall of the chamber. Finally, one gradually requires him to rear up closer and closer to the bar. If one's patience, observational skills, and timing are average, it is possible to have Sniffy barpressing frequently in 40–60 min.

Conditioning

If the student is a successful shaper, the time will come when Sniffy will press the bar three or four times within a minute. When that happens, the animal is starting to show conditioned responding on a continuous reinforcement schedule (CRF). The cumulative record in Figure 1 shows an accelerated version, in which the acquisition and extinction of barpressing takes about 20 min instead of the normal 60 min.

In the first 9 min, Sniffy presses the bar about 11 times while the student is reinforcing rearing near the back wall. Over the next several minutes, the response rate climbs; the cumulative record becomes steeper and steeper. When this happens, one can see that learning is (in part) a matter of changing the probability of occurrence of existing behaviors. Operant conditioning gives us the technology for accomplishing these changes in animals and in people.

Extinction

At this point in the experiment, if the rat is barpressing frequently, a Sniffy file should be saved for use in the second lab. The next step is to observe the phenomenon called extinction. The extinction procedure consists of stopping reinforcement. In Figure 1, the last reinforcer is delivered at about the 12-min mark. As a consequence of this procedure, the reinforced behavior will become less frequent until eventually the barpressing response that was conditioned will occur no more frequently than it did before conditioning.

To institute extinction, the Training Schedule option is selected from the Schedule menu, and the radio button labeled Extinction is toggled. This means that Sniffy will effectively never get reinforced. Over the next several minutes, Sniffy's barpressing rate will decline and the cumulative record will eventually flatten out.

LAB 2
Schedules of Reinforcement

The purpose of this lab is to place Sniffy on a PRF schedule and observe that the schedule enhances resistance to extinction. Thus far, we have talked about reinforcement as if it were something that had to occur on every occasion. Every time Sniffy pressed the bar, he got a pellet of food. This is continuous reinforcement, or CRF. One could choose, however, to deliver a reinforcer after only some of Sniffy's barpresses. Reinforcing only some instances of a behavior pattern is partial reinforcement, or PRF.

CRF is the most efficient way to shape up a new behavior quickly, but it is no longer necessary once the new behavior has been established. A judiciously chosen schedule of PRF can maintain a behavior indefinitely. PRF also has the advantage of enhancing resistance to extinction. An animal that has been partially reinforced will make many more responses in extinction than one that has been continuously reinforced.

A schedule of reinforcement is a rule for determining which responses to reinforce. There are two basic families of schedules: ratio schedules and interval schedules. Ratio schedules reinforce the subject for making some number of responses. On a fixed ratio (FR) schedule, the number of responses required is always the same. On an FR5 schedule, the subject must make five responses for each reinforcement. This is rather like being paid for piece work, with the amount of money earned determined by the amount of work accomplished according to a prearranged wage scale. On a variable ratio (VR) schedule, the value of the schedule specifies an average number of responses that must be made, but the exact number varies from reinforcement to reinforcement. On a VR5 schedule, the subject is reinforced for every five responses on the average. Las Vegas slot machines pay off on VR schedules.

Interval schedules reinforce the subject for the first response made after a specified time interval has elapsed since the last reinforcement was received. On a fixed interval (FI) schedule, the interval that must elapse before the next response is reinforced is always the same. On an FI 10-sec schedule, the next response to be reinforced will be the first response that occurs after 10 sec have elapsed following the previous reinforcement. On a variable interval (VI) schedule, the time interval following reinforcement that must elapse before the next response is reinforced varies from reinforcement to reinforcement. On a VI 10-sec schedule, the time would randomly vary from 1 to 20 sec with an average of 10 sec. Once the interval has elapsed, the reinforcer becomes available and remains available until the subject responds.

All PRF schedules enhance resistance to extinction, but the degree of enhancement depends on the kind and value of the schedule employed. In addition, each of the four types of schedules maintains a different characteristic pattern of responding when the cumulative record is inspected. However, describing and explaining these differences are topics that are usually beyond the scope of introductory psychology courses.

To train Sniffy on a PRF schedule, one starts with a Sniffy file in which Sniffy has already been trained to barpress for continuous reinforcement. If the student successfully conditioned Sniffy during the first lab and saved the file before extinction, then that file can be opened to continue with this lab. Otherwise, one can use the file called BARPRESS included with the Sniffy software package.

PRF schedules are chosen by using the Schedule menu. Ratio schedules are chosen by clicking the radio button labeled "Responses." Interval schedules are chosen by clicking the radio button labeled "Seconds." Fixed (FI and FR) schedules are chosen by clicking the ratio button labeled "Fixed," and variable (VI and VR) schedules are chosen by clicking the button labeled "Variable." Schedule values in responses or seconds are specified by typing an integer in the number entry box.

Sniffy will extinguish if one selects too large a value when he is first placed on a PRF schedule. Real animals can be induced to continue responding on PRF schedules where the amount of energy they expend responding is greater than the amount of energy they can derive from the reinforcers. However, to get them to do so, students must start out with small response or time values, increase the values gradually, and allow the animal's behavior to stabilize at each value before moving on to the next. The cumulative record will show high response rates as FR10, FR20, and FR30 schedules are acquired with the typical "pause and run" pattern emerging on the largest schedule (see Figure 2 below).

All students should start with the same file (or their own rat from the previous lab) where Sniffy has been on CRF for at least 10 min, maintaining a steady rate of responding. The emphasis of this lab is to require students to document every change in the schedule they initiate and record the effects. There are three phases, all of which may require students to measure the time it takes for extinction to occur. For this purpose, we recommend producing a "ruler" that they can hold up to their cumulative record on the computer screen.

The three phases are summarized by the following three questions students are required to answer during this lab. (1) How long does it take for barpressing to extinguish after a rat has been trained on a CRF schedule? (2) When changing from a CRF schedule to a PRF schedule, what is the largest value that will maintain barpressing? (3) What is the largest value on a PRF schedule that you can train Sniffy to maintain barpressing on? The instructor may want to assign FR, FI, VR, and VI schedules to different groups of students, since there is time only to explore one schedule in detail.

Students need to be clear that while an animal may not learn a VI50 schedule directly after CRF training, the animal can be shaped through successive schedules (e.g., CRF to VI20, VI20 to VI40, VI40 to VI80, . . . etc.) to eventually maintain barpressing on schedules well over 100. This step requires a lot of patience, as well as a trial and error approach. Students will spend some time staring at Sniffy while waiting to determine the outcome.

The students are required to document every step, reporting whether the behavior extinguishes, in which case they measure how long it took to extinguish since the last schedule change, or whether the behavior was maintained, in which case they are asked to save the animal under the FILE menu. It is always a good idea to save a Sniffy file before each increment, particularly if one decides to get adventurous and try a larger than average step increase. One of the big advantages Sniffy has over a real rat is that if Sniffy files are saved regularly, one will not have to recondition him from scratch if he does extinguish. The assignment for this lab could be to summarize the procedures and results, hand in a cumulative record of the finished product, and answer take-home questions.

We conduct student evaluations twice during the labs to quantify students' preferences among the eight classes of software that we employ during the year. Sniffy was ranked second overall, even though students claimed it was one of the harder programs to learn how to operate. Part of this preference must be due to the interactive nature of the task. The students can see how their carefully timed behaviors begin to affect the rat's behaviors in the Skinner box.

The shaping lab was clearly more interesting to students than the PRF lab, primarily, we speculate, because the second lab was much more passive. Students change the schedule and spend many minutes waiting to see whether barpressing is extinguished or maintained (as they would with real rats as well!). We have tried to make the PRF lab more interesting by having students work in groups of eight who plan the use of their four computers to answer the three main questions for each of the four schedules. Such collaboration seems to work well, even though it is nearly always the case that two or three students do most of the planning (and delegating).

We have recently developed a behavior encoding device that allows students to press keys on the keyboard assigned to the 10 major classes of behaviors that Sniffy (or any other organism) emits to score baseline observations. This program then dumps the results to a central server that computes interrater reliability scores. Thus, we can also include in our curriculum important issues about observational techniques.

Sniffy can also be used to discuss artificial intelligence and simulation issues in cognitive science. The next section describes some of the algorithms that we employ to generate realistic random behaviors that gradually come under the control of reinforcement contingencies. An advanced course could study these algorithms as a model of real operant learning. This would introduce the distinction between hard and soft AI and provide exposure to modeling and performance evaluation techniques.

THE PROGRAM

Animating Sniffy's Movements and Actions

The animation and learning routines were the most difficult components of this C program. The animation was accomplished by starting with video clips of a real rat wandering around a terrarium. Approximately 15 basic actions were selected for Sniffy's repertoire, including sniffing, walking, turning, scratching, rearing, drinking, eating, and genital licking. All of these categories had at least two versions (normal and mirror image), and some had more (e.g., walking north, east, south, and west, NE, NW, SE, and SW). A medical artist converted each of these video clips to a series of PICT files, which when played at 5-10 frames per second would animate that action sequence. Each one of these behaviors is called a *sequence*.

When a sequence is played, the frames within the sequence are displayed sequentially as a cartoon animation. At the end of the sequence, the program determines which sequence to play next. This is done randomly, with constraints imposed so that Sniffy stays within the chamber and the transitions between sequences are relatively seamless and smooth. At the outset, each sequence has a relatively small probability of being selected (a modifiable parameter set to default values in the resource file), and depending on where the animal ends up in the chamber, only a subset of the available moves are legal.

Sniffy is able to learn (and forget) actions (or sequences) and locations (called *sectors*) that form associations with the reward. The program does this by pairing the number of occurrences of the reward with the location within the Skinner box and with the sequence that has been performed. Figure 2 shows the sectors of the chamber seen with the programmer's debugging window overlaid on the Skinner box window. There are 9 floor sectors and 10 wall sectors that can become *attractors* if food is provided consistently while the animal is at that location.

Location learning is functional only if the rat is "close enough" to the hopper, represented by a circular crease around the food hopper on the debugging window. "Close enough" expands as the rat becomes magazine trained. If the reinforcement is presented when the rat is within the "close enough" circle, the circle is expanded and the rat moves toward it. Likewise, it is reduced if the rat was outside the circle when the reinforcement was presented. When the circle has expanded so that the whole Skinner box is within the circle, the rat will always move to the food and so is said to be magazine trained.

Selection of the next sequence of animation to play is determined by a routine that determines what sector the rat is in and whether conditions are right for the rat to perform any actions (e.g., it can eat if it is at the hopper and cheese is in the hopper). The routine then makes two passes through the list of all sequences. On the first pass, each sequence is considered, and a determination is made as to whether the sequence can be played under the present conditions. This determination also returns the number of associations of the attracting sector if (and only if) the sequence moves the rat "closer" to that sector. For each of the valid sequences found, the base probability (or frequency of occurrence) is added to a running total. The frequency is adjusted, depending on a number of conditions, in order to ensure that schedule effects come out right. A random number is then chosen between 0 and the total frequency of all the playable sequences. This number is used to select the next sequence in a manner that preserves the sequences' relative probablilites.

Learning on a Logistic Curve

Once the rat is magazine trained, it can be shaped. This shaping influences the activities it performs (sequences) and the locations it frequents (sectors). The shaping of sequences and sectors are only partially independent of each other. Each is based on the number of associations. This value will increase with pairing. Both can also decrease, but a sector's associations will decrease only if a sequence is decreased as well. The following sections describe sequence and sector shaping in more detail.

Each sequence carries three variables that are related to shaping. Base_freq is the relative probability of this sequence for an untrained rat—that is, the value that the rat starts with. The variable *associations* is a counter of the number of pairings of this sequence with the reinforcer. The function $f(x)$ is a member of a family of curves called the *logistic function* used in some neural network

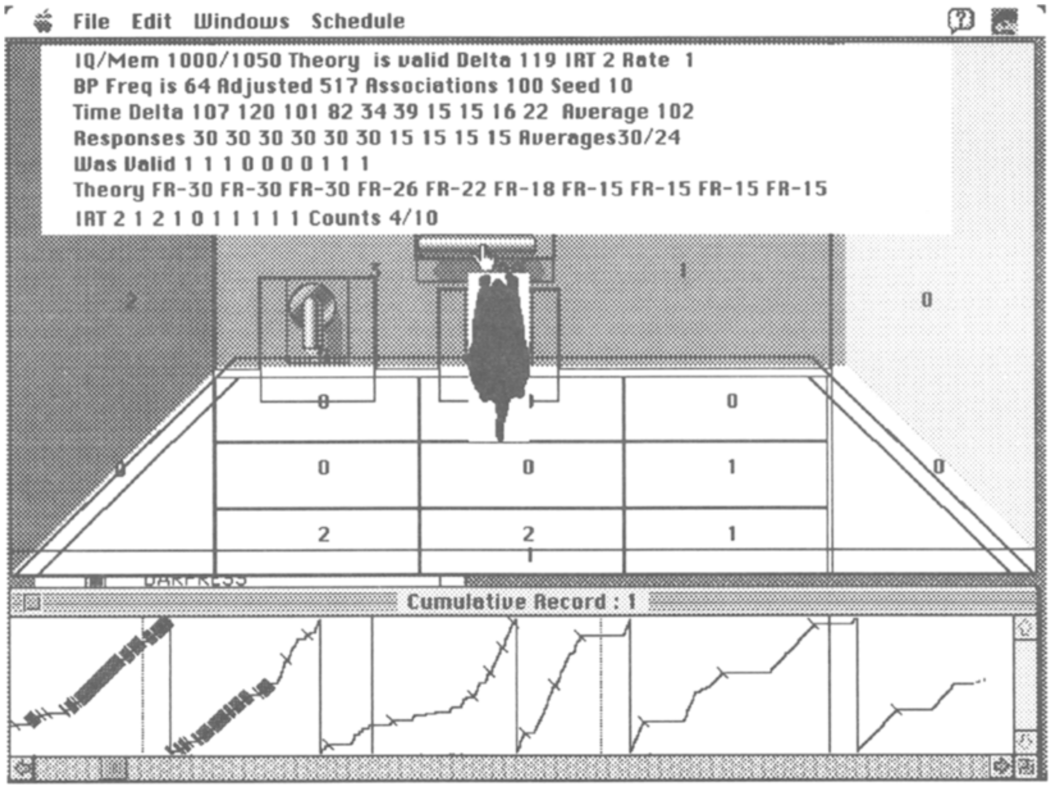

Figure 2. Debugging screen superimposed over the operant chamber showing the learning and memory parameters during FR 15 training, and association frequencies in location sectors.

learning algorithms and displayed in Figure 3. The base_freq is added to *f*(associations) to produce the value *freq*. This value, shown on the *y*-axis, is then used as the relative probability of this sequence during selection as a function of the number of reinforcements (i.e., associations) shown on the *x*-axis.

The number of associations controls the amount of learning. As mentioned before, associations increase and decrease. Increases take place in two ways. In the simplest case, associations (both sector and sequence) are increased by one when the rat is presented food. Other cases are more complicated and are therefore beyond the scope of this paper. Any time the rat performs a sequence and does not get food, that sequence's probability is a candidate for decrementing. Naturally, if the sequence has no associations, no decrementing is performed. Location sector associations are decremented by the same amount in most cases.

Both the sequence and the sector have a maximum value established beyond which associations will not increment. This prevents the rat from accumulating "too many" associations and ensures extinction in a fixed, short time. Without the maximum, the associations would continue to build. If the rat were left on an FR40 for several hours and then put on an extinction schedule, it would take several hours for the behavior to extinguish. Because the maximum is relatively low, extinction in this case would occur in the same time (5–10 min) as it would for a rat that had been on FR40 for only 10 min.

Determining PRF Schedule Effects

The program maintains three variables that correspond to the schedule that the rat "thinks" it is on. Guess_Responses, if true, means that the rat thinks that it is on ratio schedule; false implies an interval schedule. Guess_value is the size of the schedule. Guess_fixed is true if the rat is responding as if the schedule was fixed; false implies a variable schedule. An additional variable (theory_valid) is set to false if the rat discovers evidence that its current guess at a schedule is wrong.

Initially, the rat supposes a CRF schedule. Each time the rat's behavior is reinforced with the food, the rat "remembers" the occurrence. With arrays of Size 10, the simulation "remembers" the last 10 occurrences and records the delta time since the last reinforcement, the number of times Sniffy has reared in the sector with the most number of associations, and the four variables corresponding to the guess at the current schedule.

Sniffy continues to think his theory is correct until he gets evidence otherwise. This happens when he either does not get a reward (when he has expected one) or gets a reward when he has not expected one. The expectation

```
y = a/(1+e(-x + b)/c )
```

The default parameter values used are:

```
a = max_freq_factor = 60
b = behaviour_threshold = 40
c = scale_slope_factor = 8
```

This yields the curve :

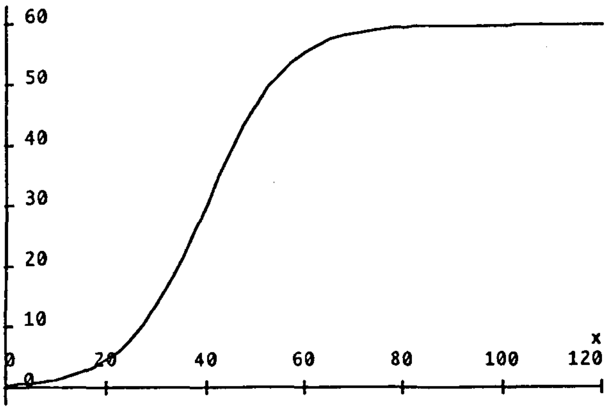

Figure 3. Default logistics function used in the learning algorithms, showing the relative probability of a behavior as a function of the reinforcement frequency.

is based on his guess value for the schedule, so if Sniffy thinks that the schedule is FI20 and presses the bar 25 sec after the last reinforcement, he will think that his theory is incorrect. Similarly, if he presses after only 15 sec and gets a reward, Sniffy will think that his theory is incorrect.

When the rat thinks that its theory is incorrect, it tries to construct a new theory. To do this, the rat determines the schedule type for the current theory. Two instances allow the rat to switch back to a fixed schedule from a variable one. The rat switches from VR to FR if each reinforcement has come after the same number of reinforcements. The rat switches from VI to FI if no reinforcement has come with a delta time less than the current delta time.

Modifying the Simulation Parameters

The program allows the user to modify learning parameters and to customize the display to suit a variety of Macintosh platforms. These parameters control several different aspects of the simulation. To aid this process, a programmer's debugging screen is overlaid on the Skinner box window when option-shift is held down while a password is typed. In Figure 2, the debug screen shows some of the learning variables affected by parameter manipulations, as well as the memory vectors that Sniffy uses to "figure out" what schedule of reinforcement seems to be in effect.

The sophisticated user can modify the behavior of the simulation relating to the computer environment, how quickly the animal trains, how and when schedule effects become apparent, and how pronounced they are. To change any of these parameters, the advanced user employs ResEdit (available with MacLaboratory or from Apple dealers) to open the resource ID that needs changing. The Sniffy software does very little error checking on the values of these parameters, so one must be sure to test any changes thoroughly.

Availability and Future of Sniffy

Sniffy is available from MacLaboratory, Inc., 314 Exeter Rd., Devon, PA 19333, for $49.95 per CPU up to 10 units and at half price for additional units. We hope that Sniffy's future will be bright. We have priced the product very reasonably in order to recoup some of our development costs to reinvest in an enhanced Version 5.0. There are many phenomena in operant paradigms that could be incorporated. Some of these include the effects of satiation, spontaneous recovery, punishment, and negative reinforcement. More ambitious improvements would require rethinking the learning algorithms and providing more sophisticated ways of representing "environmental knowledge."

For example, to simulate discrimination learning, we would have to provide a way of encoding the stimuli in the presence of which the reinforcement contingencies hold. For matching-to-sample experiments, we would have to implement a very different sort of forgetting that was more time based (as would be required to exhibit spontaneous recovery). Our early thoughts are leading us to consider neural network learning algorithms and distributed representations that may give us the power we need. Such a development would extend the usefulness of Sniffy in the classroom as a toy problem domain that we could employ to teach connectionist simulation techniques.