Learning Movements by Association without Reinforcement: Edwin Guthrie's Contiguity Theory Revisited

Clayton Lewis
University of Colorado Boulder
December 23, 2021

Abstract: In the decades 1920-1960, a period of extended controversy in the theory of learning, Edwin Guthrie argued that in tasks like maze learning, or escape from a puzzle box, animals learn to act by forming appropriate associations between stimuli and responses in a single trial, with no important role for reinforcement. While some of Guthrie's ideas, especially as developed by W K Estes, have since come to dominate many areas of learning theory, his account of learning sequences of movements has largely been forgotten, in favor of accounts that give reinforcement a central role, and see learning as a graded process. But can Guthrie have been right, after all? This paper presents modeling results suggesting that Guthrie's ideas can account for some aspects of movement learning. At the same time, some of Guthrie's thinking seems to be off target. Apart from their possible theoretical interest, the results cast light on the way in which ideas are formed in intellectual controversy in psychology.

The study of the behavior of animals in mazes has long been important in psychology, with the first studies being reported in Small (1901), and work continuing to the present (see Thomas et al., 2014, for a historical review, and Benvenuti et al., 2021, and Rosenberg, et al. 2021, as current examples.) Of less current interest are studies in which animals learn to escape from a cage, but these studies, too, have been important in the development of psychological theory, starting with reports by Thorndike (1898). The basic phenomenon of interest in most studies is that animals navigate a maze, or escape from a puzzle box, faster and faster, over repeated trials, learning to enact sequences of movements that, over time, become more efficient and direct. How does this happen?

The purpose of this paper is to revive and reexamine a particular set of ideas on this question, those of Edwin Guthrie (1886-1959). From about 1930 to 1950 Guthrie was one of the most prominent American psychologists (for example, he served as president of the American Psychological Association in 1945; see Sheffield, 1959), and a very active controversialist in learning theory.

A central investigation in Guthrie's career was an extensive study of cats escaping from puzzle boxes, that is, cages that contained some kind of trigger that would open a door (Guthrie and Horton, 1946). In accounting for the behavior of the cats, Guthrie claimed that "stimulus patterns active when an act is initiated will serve on later occasions to favor the reinstatement of the act. (Guthrie, 1940)" Further, Guthrie claimed that the associations between stimulus patterns and movements are formed all at once, and not gradually, over multiple trials: "The position taken in this paper is that the animal learns to escape with its first escape (Guthrie, 1940)."

Guthrie's ideas came to be known as "contiguity theory", because of the claim that stimuli and responses are associated simply by occurring together. The simplicity of the assertions, together with Guthrie's claim that these were the *only* ideas really needed to understand learning, brought Guthrie into conflict, in one way or another, with most of his contemporaries. In particular, he dispensed with a role for reward or reinforcement, and this separated him from many. In Guthrie's view, the effect of a reward is simply to change the stimulus situation, so as to protect already-learned associations from interference: "All that escape or reward does is to protect the learning from being unlearned (Guthrie, 1939)."

The effort to make these simple ideas work led to the development of stimulus sampling theory by W.K. Estes (1950; see anticipation in Voeks, 1948). By decomposing "a stimulus" into a collection of many stimulus elements, Estes constructed a model in which one gets the effect of graded association strengths, when fewer or more stimulus elements have been associated with a response, even though the individual associations are all or nothing. As Bower (1994) describes, Estes's work led to a flowering of models of learning that survives to the present, including connectionist models.

Thus one can say that, in the long run, Guthrie came out well from the controversies of his time. Contemporary models of memory owe more to Guthrie than to Ernest Hilgard, B.F. Skinner, or Edward Tolman, who were among those with whom he argued (Guthrie, 1935, 1940). However, in the domain of maze learning, Guthrie's ideas have been largely left behind. Also, as Bower recounts, while Estes tried to follow Guthrie in his limited view of the role of reinforcement, he (and others) abandoned the effort. As Jensen (2006) documents, current discussions of maze learning in psychology textbooks are dominated by ideas from Edward Tolman's work, centering on the ideas of "latent learning", a form of learning that takes place without rewards being provided, and "cognitive maps", representations of the geometry of the maze that the animal is supposed to construct. As just mentioned, Tolman was another of Guthrie's disputants, and in this area, at least, Guthrie's ideas have not survived.

As Jensen notes, however, the extended controversies over maze learning were never really settled. In the case of latent learning, "The end of an era of inspired research productivity and sharp theoretical debate had come to an end not because one theory prevailed over all others. Rather ... it was because the issues that arose out of the extensive latent learning experimentation remained unsettled, and no resolution was thought to be forthcoming. Thus, what had begun as a lively, empirically based debate over fundamental issues in learning ultimately ended in a stalemate (Jensen, 2006)" Against this background we can still ask, could or could not Guthrie's simple ideas provide a workable account of maze learning? In addressing this question we can use a tool not available to Guthrie and his co-controversialists, computational modeling.

As mentioned earlier, Guthrie focussed much of his attention on puzzle box escapes, rather than mazes, but the two tasks are essentially the same, as Guthrie views them: the animal needs to learn a sequence of movements. Accordingly, we seek to develop a model animal (a cat, let us say) that instantiates Guthrie's key assumptions, and see if it can produce appropriate learning, in puzzle box escapes and in mazes.

Our model will instantiate Guthrie's ideas in an extremely simple way. The "stimulus situation" at any time during a trial will be identified with the cat's position in a two-dimensional grid of squares. The cat will have just four possible movements in any situation, moving north, east, south, or west. These movements will take the cat into a new stimulus situation, unless the destination square is occupied by a barrier or obstacle.

When the cat makes a move, that movement is associated with the stimulus situation in which the movement is made, that is, with the square the cat is in when it makes the move. When the cat enters a square it has visited before, it makes the same move it made last time. When it enters a square it has not visited, it randomly chooses a permissible move. One of the squares is the escape trigger: when the cat enters that square it is deemed to have escaped, and the trial ends.

Note that the model follows Guthrie's claim that associations between stimulus situations and movements are all or none, and not graded. Note also that there is no reward or reinforcement mechanism in the model, as Guthrie would have had it.

It might seem that even this simple model has some chance of working. But when it is run, a fatal problem immediately appears: the cat almost always goes into a loop, and never escapes. This issue seems never to have been fully appreciated, by Guthrie or anyone commenting on his ideas, despite Guthrie's stress on the idea of repeating previous movements. Guthrie and Horton (1946) note that some actions are self-terminating, that is, cannot be repeated, but movements in a cage are not in that category. Indeed, Guthrie's whole analysis rests on the premise that situations in the cage, and movements, *are* repeated across trials. Some means is needed to prevent actions being repeated fruitlessly within trials.

Once the necessity of breaking loops in the cat's behavior is recognized, various possibilities can be explored. Perhaps the simplest is introducing noise: when a previously-visited square is reached, with some probability, the move made previously is not made, but a random legal move is chosen instead. In this scheme any loop will be broken eventually, since there is some finite chance of any given move being replaced.

Figure 1 shows results for a model cat placed in square (0,0) in a cage 20 squares on a side, with the escape square at (17,17). The probability of making a random move, when a stored move is available, is .05. Panel (a) shows the cat's trajectory on the first trial, and Panel (b) on the fifth trial. Because the cat often visits the same cells many times, a move into a given cell is drawn to a different location within the cell, on successive visits, so as to make it a bit easier to follow the sequence of repeated moves. Early visits are drawn to points near the upper left corner of a square, and subsequent ones are drawn farther to the lower right. Panel (c) shows the number of moves to escape on 20 trials.

Figure 1a. Trajectory of a cat on its first escape. The cat starts in the NW corner of the cage. About three-quarters of squares in the cage are visited several times. After exploring the NW, NW, and SW quadrants of the cage quite extensively, the cat eventually moves into the SE quadrant, where the escape square is located, and, after dithering around for a while in that area, it eventually escapes. About a quarter of the cage has not been visited at all. See text for explanation of how the moves are drawn in the trace.

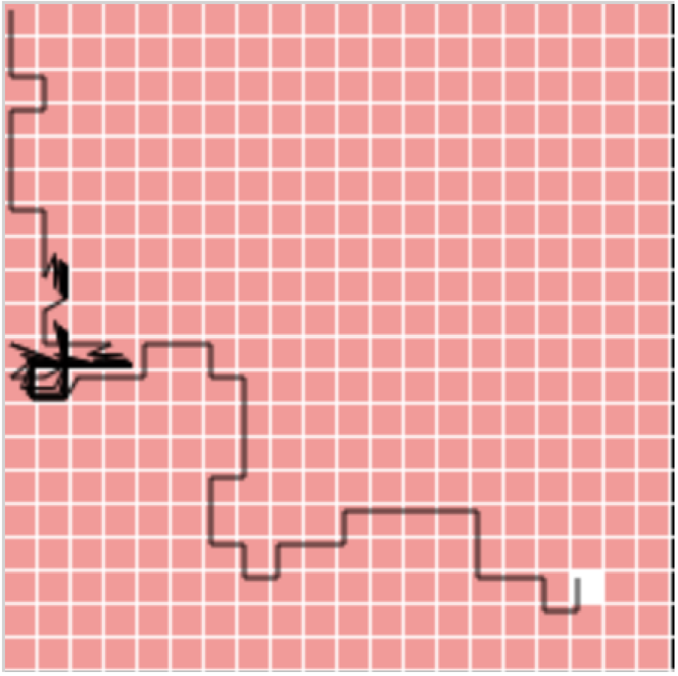

Figure 1b. Trajectory of a cat on its fifth escape. After a bit of dithering in the NW quadrant, the cat moves east, and then down, and then east again, reaching the escape square with no further looping. The total amount of looping is far less than on the first trial.

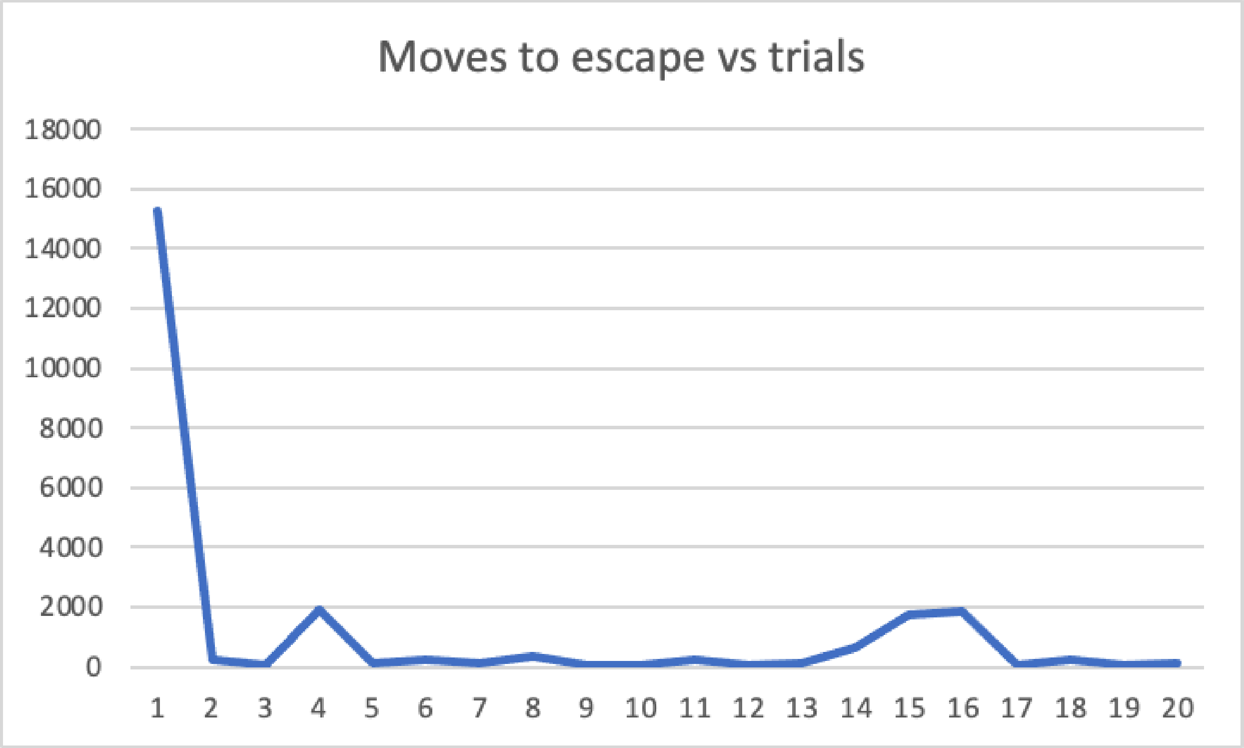

Figure 1c: Moves to escape plotted against trials. The first escape takes about 15,000 moves, while none of the subsequent escapes take as many as 2,000 moves, with most taking less than 300 moves. The shortest possible escape takes 34 moves, but the cat never takes less than 42 moves in this series.

As can be seen, the model cat does escape much more quickly on later trials than on the first one. That this is generally so, and not a chance occurrence for this single run, can be seen in Figure 2, that shows the median moves to escape for 50 model cats making 20 escapes each: the median falls from about 19,000 moves on the first trial to about 250 on the 20th.

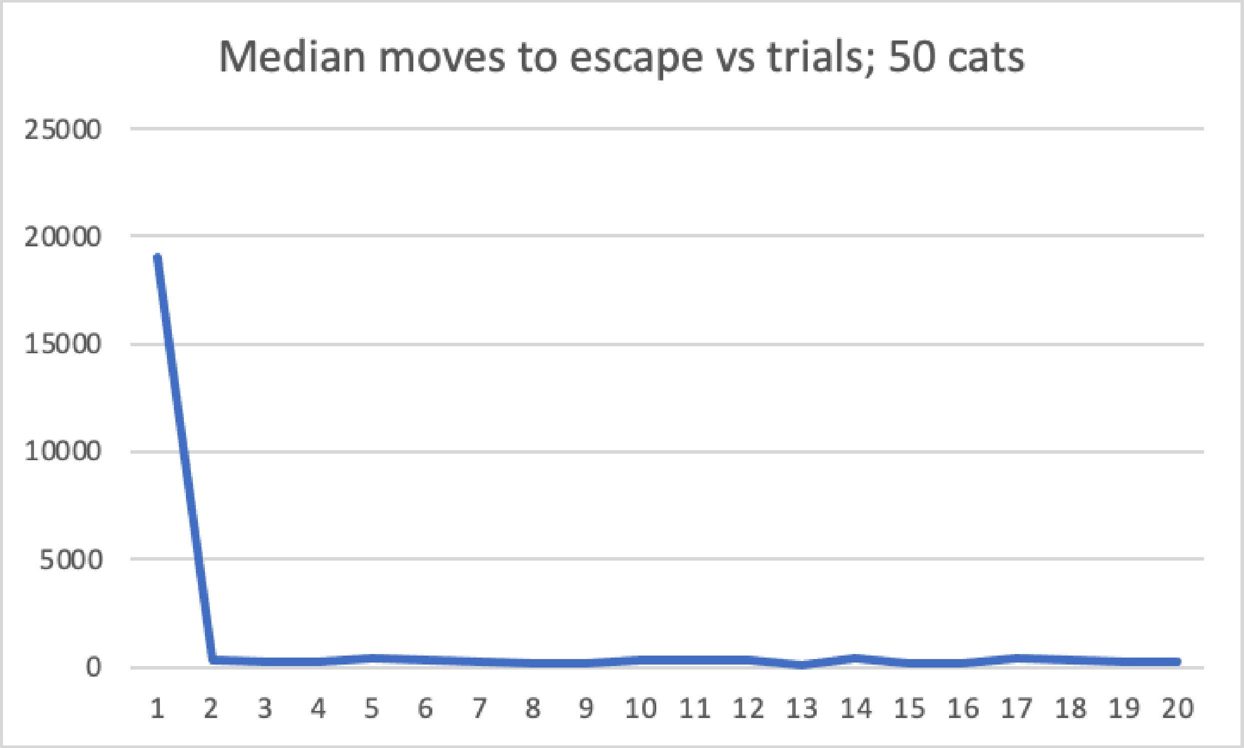

Figure 3: Median moves to escape, for 50 model cats. The median falls from about 19,000 on the first trial, to between 125 and 435 for subsequent trials.

No reward figures in the model, yet performance improves dramatically. Does this improvement vindicate Guthrie's ideas? In a way, yes. We do indeed see that reward is not necessary for learning to occur in this situation. And Figure 2 shows that essentially all of the improvement happens on the first trial, just as Guthrie and Horton(1946) argued: "The cat learns to escape in one trial. (p41)". How does this happen?

Some insight into the workings of the model can be gained by considering the information that the model stores as the cat moves through the cage. Each time the cat moves from one square to another, the move it makes is stored in the square from which it was made. This leaves what we can call a *trace* in the squares: for each square the cat has visited, the trace contains the move the cat last made from that square. On each trial after the first, the cat follows the trace, except for any noise moves that may be introduced.

Why are the second and following trials shorter than the first? One might think that the second trial, following the trace of the first, would be just the same as the first. But that overlooks the *loops* that are virtually always going to be found on the first trial. The trace of any trial, however many loops the trial contains, contains *no* loops, as long as the trial has led to an eventual escape (see proof in Appendix A). If there were no noise, the second and subsequent trials, being driven by the trace of earlier trials, would contain no loops at all. Thus the first trial takes a long time, because of loops, and subsequent trials take less time, because the only loops that can occur are triggered by noise, which is relatively rare.

Thus the improvement seen in this version of the model consists only of the elimination of loops. One can see this by examining not the total number of moves taken in a trial, but the number of unique squares the cat visits; in this measure, squares visited repeatedly in a loop are counted only once. As Figure 3 shows, there is no improvement at all in this measure, after the first of 20 trials.

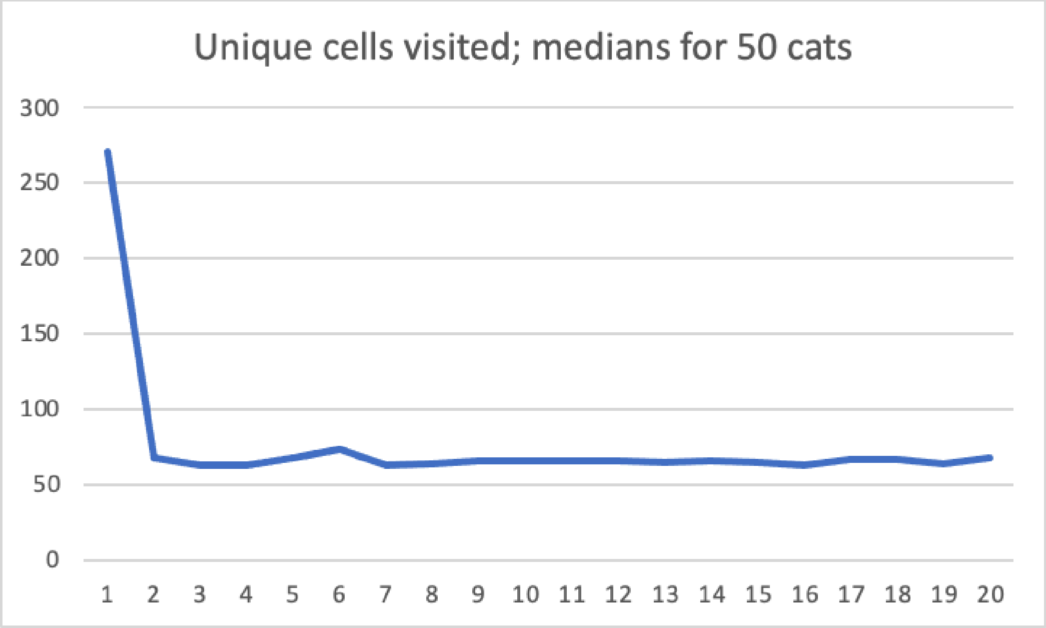

Figure 3: Median number of unique squares visited on 20 escapes by 50 model cats. The value falls from about 270 on the first escape, to about 65 for all subsequent escapes.

Before exploring further the adequacy of this model, let us consider an aspect of its behavior other than number of moves. Guthrie himself was not very interested in how long it took his cats to escape. In Guthrie and Horton (1946), in which they present the results of his studies of escapes, they say, “The total time elapsed was of minor interest to us. (p16)”. Indeed, they do not report times to escape for their study.

Instead, Guthrie and Horton focused on patterns of repeated actions. They saw these patterns as providing direct evidence for his assertion that animals simply repeat previous actions. The presentation in Guthrie and Horton includes drawings, made from cine frames, that show the pose of cats at the moment when they activate the escape trigger, for repeated escapes. It is indeed striking that many of these poses are very similar, for repeated escapes for a given cat, while being quite different for different cats. Further, Guthrie and Horton assert “a tendency on the part of the final series of movements, the action of the last few seconds... to be more stable than preliminary behavior. (p39)”

Guthrie and Horton take this as support for Guthrie’s idea that escape acts to protect actions taken just before escape from interference, that is, by being overlaid by different responses made later, when in the same situation. After the cat escapes, the argument goes, it can no longer be in any situation in the cage, and hence the last actions it took when in the cage cannot be overlaid by others. “The reason for the remarkable preservation of the end action leading to escape is that this action removes the cat from the situation and hence allows no new responses to become attached to the puzzle-box situation (Guthrie and Horton, 1946, p39).”

Guthrie and Horton did not provide any systematic comparison of repeated action sequences at different stages in escape trials, but we can easily do this, for our simple model. Of course our model has only a very spare representation of the actual actions of cats, being restricted only to moves between squares, and not including head movements, particular leg movements, and so on (see Appendix B for more discussion of this matter). Nevertheless we can ask whether the sequence of the last five actions (say) before escape are more likely to be seen in multiple trials than sequences of five actions seen earlier in an escape.

Figure 4 compares in how many trials, in a run of 30 trials, sequences of five actions are repeated exactly. The sequences compared are the last five actions before escape, the five actions ending 20 actions before escape, and the five actions ending 40 actions before escape, all as observed on the last trial. The results shown are the median numbers of earlier trials in which the sequence was seen, taken over 100 model cats. As can be seen, the final sequence is actually somewhat *less* likely to have occurred on earlier trials than earlier sequences, the opposite of Guthrie and Horton’s expectation. The number of repetitions is larger for the earlier sequence than for the last sequence for 63 of the 100 model cats. The 95% confidence interval for this proportion extends from about .53 to about .72.

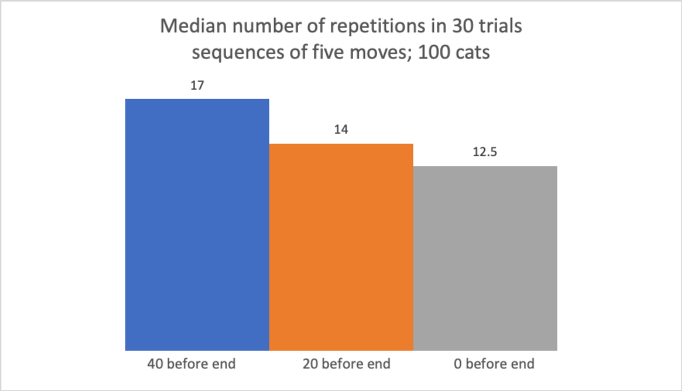

Figure 4: Median number of trials on which a given sequence of moves was repeated, in a series of 30 trials, for 100 model cats. The sequences examined are the five moves ending 40 moves from the end of the last trial, 20 moves from the end, and the last five moves on that trial.

The median number of repetitions is largest (17) for the earliest of the three segments, and smallest for the latest segment (about 13).

Why might this be? First, we can note that there is a logical flaw in Guthrie's release from interference argument, in the context of repeated escapes. When a cat escapes, it is only *temporarily* removed from the cage. Except for the very last trial, it is always put back in the cage for the next trial. Thus exposure to interference is almost the same for any actions in the cage, however near escape they are. Second, the reason actions may *not* be repeated, in the model, is noise. The later a sequence of actions is in a trial, the more likely it is that a noise move will have modified the trajectory before the sequence is reached, so that the sequence is not repeated.

Returning to moves to escape, while Guthrie and Horton were not concerned with this, other researchers were. Indeed, Thorndike's original observations include the fact that his cats became more and more efficient in escaping, eventually being able to escape right away. A further look at Figure 2, and some attention to the geometry of the cage, shows our model cats are far from accomplishing that. In all cases, the shortest path to escape is 34 moves, but the figure shows that the median path length for the model cats, even on the 20th trial, is about 250 moves. So the model cats are simply not finding a shortest escape.

Guthrie does not address this matter. Is there a way to do so, while keeping the spirit of Guthrie's thinking? A second model attempts to do this. It includes a way to shorten escape paths, while also providing a way to break loops, without resorting to noise, as the first model does.

In this second model, which we will call the *path reduction* model (PR, for short) the basic trace mechanism is unchanged, reflecting Guthrie's commitment to animals repeating actions when revisiting situations. Two operations are added. First, the model recognizes one simple form of loop. Suppose the cat finds that it has moved into a square in which there is a stored move that would take the cat back to the square it has just left. For example, suppose square S0 contains the move E, for east, and the square to the east, S1, contains the move W. This situation traps the cat in a loop: starting in S0, the cat makes the move E, and reaches S1. There it finds the move W, that directs it back to S0, and so on. To avoid this, when the model cat moves E into S1, the W stored in S1 is replaced by a random permissible move. This random move could be W, in which case the loop is not immediately broken. But eventually one of the moves in S0 or S1 will be replaced by a different move, and the loop is broken. We can call this operation BackupRemoval (though it has that effect only probabilistically.)

While this is not an idea contemplated by Guthrie, it stays within the broad framework of his thought. There is no idea of reward, nor is any global information, about the layout of the cage, or the overall geometry of the cage, or the cat's trajectory, such as Tolman posited, captured or used. The stored move is adjusted in response to a situation that can be identified using purely local information, as the cat moves in the cage.

A second local adjustment operates when the model cat finds that it is turning a corner. For example, suppose square s0 contains the move E, and the square to the east, s1, contains the move S. When the cat arrives at s1, its previous move was E, and it is about to move S. Let us give the name s2 to the square that lies south of s0. If s2 is open, that is, if it is not a barrier square into which the cat cannot move, the cat can *flip* the corner, by replacing the move stored in s0 by S, and storing the move E in s2. The cat moves into the square that lies south of s1, say s3, as it would have without the flip. But now the trace would direct the cat from s0 south to s2, and then east to s3, rather than east from s0 to S1, and south to s3. Thus the corner ES has been flipped to SE, and the trace no longer visits s1. See Diagram 1.

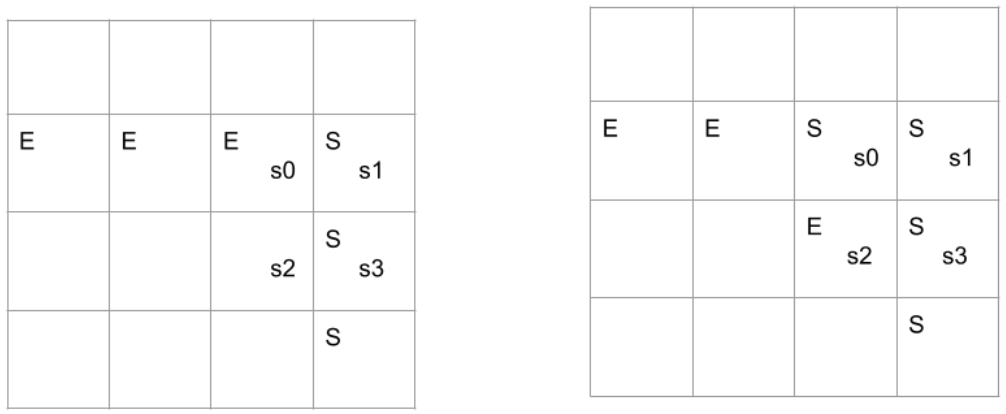

Before flip

After flip

Diagram 1. Compares a trace before and after a flip operation. Before the flip, the trace represents moves E to square s0 and then s1, and then a move S to s3. After the flip, the trace represents moves E to s0, S to s2, and E to s3. In both cases a cat following the trace reaches s3, but before the flip it moves via s1, and after the flip it moves via s2. Note that if s2 contains an obstacle, the flip cannot be carried out.

Like BackupRemoval, the Flip operation uses only local information. The information used is readily available to the model cat as it moves in the cage, and the adjustments to the trace are limited to the immediate vicinity.

When used together, the two operations RemoveBackup and Flip have important effects, as shown in Figure 5. As in Figure 1, Panel (a) shows a cat’s trajectory on its first escape, Panel (b) shows the trajectory on its fifth escape, and Panel (c) shows move counts for 20 trials. The added panel (d) shows the trajectory on the 20th trial. In these trials RemoveBackup and Flip are applied probabilistically with each being applied on 90% of the occasions when it would be applicable. Figure 6 shows the median moves to escape for 50 cats making 20 escapes each.

Figure 5a: Trajectory for first escape of PR model cat, rendered as in Figure 1. The cat has explored almost all of the squares in the cage, many times.

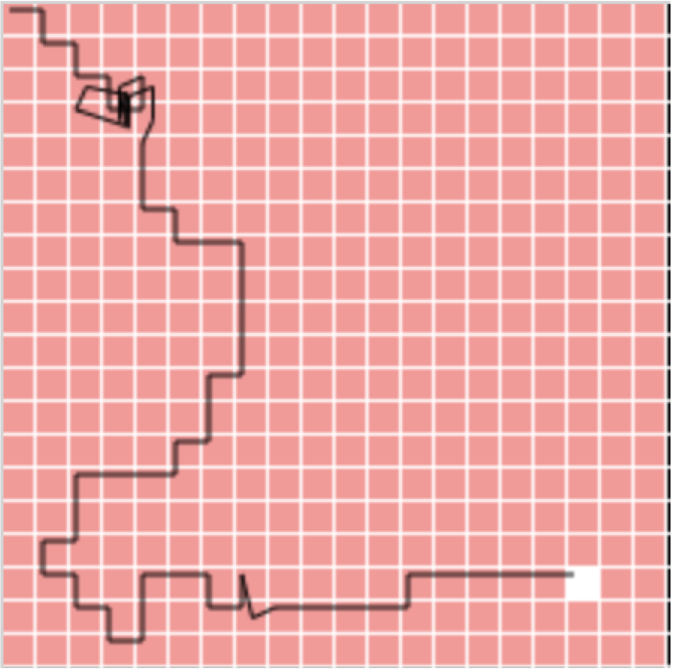

Figure 5b: Trajectory for fifth escape. After a little dithering in the NW quadrant, the cat makes only one repeat visit to any square, and visits very few squares in the rest of the cage.

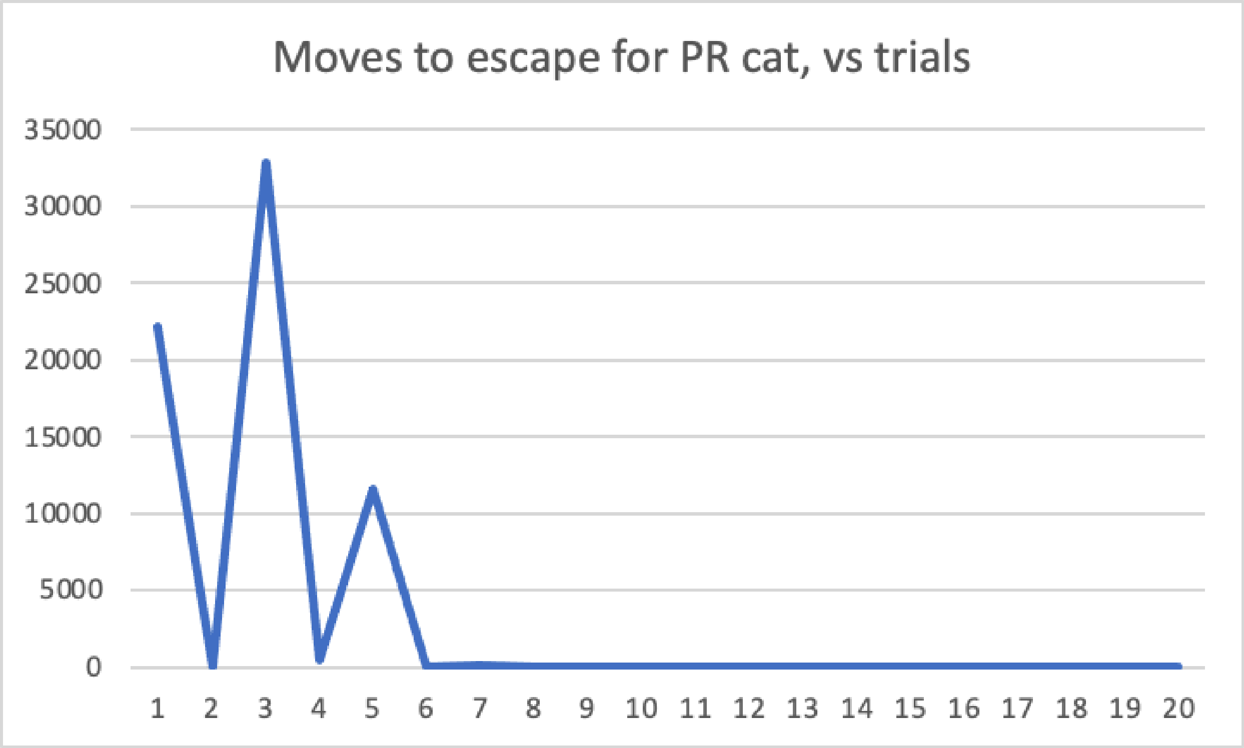

Figure 5c: Moves for 20 escapes for a PR model cat. After thousands of moves on some of the first five trials, the cat escapes very quickly thereafter. The curve is quite jagged over the first half dozen trials, but then settles down to efficient escape; both of these features are present in Thorndike's data.

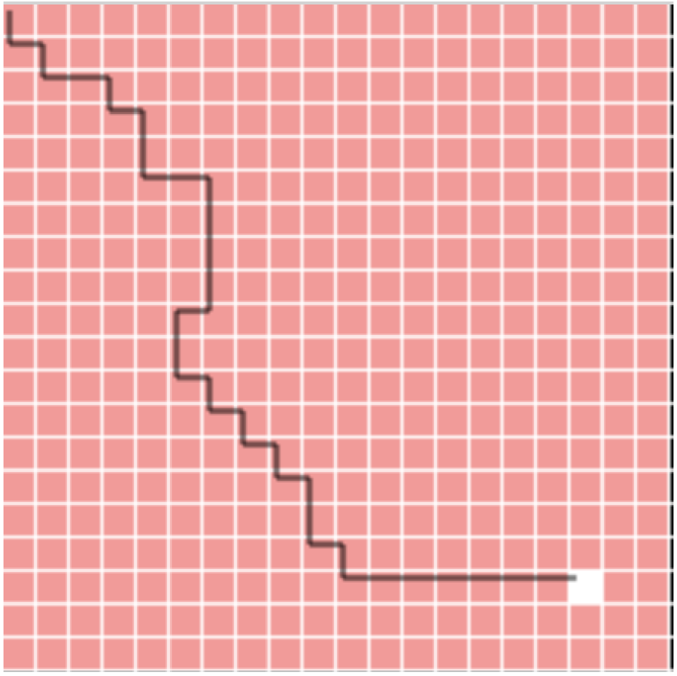

Figure 5d. Trajectory of a PR model cat on its 20th escape. It revisits no squares, and makes only two moves more than the minimum required to escape, making one unnecessary westward move, on about its 15th move, that has to be canceled by an eastward move.

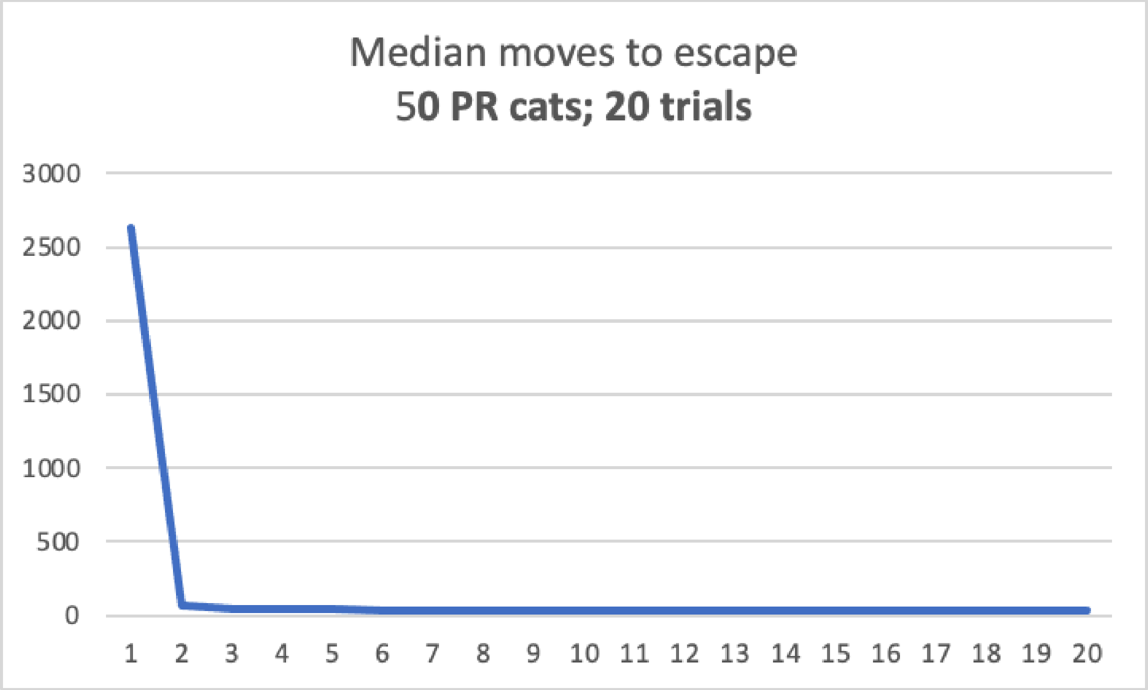

Figure 6. Median moves to escape, for 50 PR model cats. After more than 2500 moves on the first escape, the median falls to into the 60's on the second trial, the 30's by the sixth trial, and reaches the minimum possible, 34, for 11th and subsequent escapes.

Note that loop breaking has been achieved, as shown by the fact that all of the cats escape on all trials (a cat caught in a loop can never escape.) In contrast to the first model, no noise has to be introduced. We'll discuss how this breaking of loops is achieved, shortly.

Note also that, in contrast to the first model, the cats do approach the minimal length of escape, 34, as they make repeated escapes. This is what we would expect from Thorndike's observations.

Both of these desirable characteristics of the model stem from key properties of the Flip and RemoveBackup operations, as they are applied over a cat's trials. Informally, these properties are (1) that any minimal escape can only be transformed into another minimal escape, by the action of these operations, and (2) that no non-minimal escape, or trace that does not lead to an escape, is stable under these operations, except under special conditions that cannot hold in the cage we are using. Together, these properties mean that loops are eventually broken, since they are not stable, and traces that lead to escape, but are not of minimal length, are eventually transformed into paths of minimal length.

Consider property (1) first. In our cage, the starting square is (0,0), and the escape square is (17,17). There are many shortest paths connecting these, such as making 17 moves east, followed by 17 moves south, or making 17 moves south followed by 17 moves east. In fact, any sequence that includes 17 moves east, and 17 moves south, will be a path of minimal length. and all minimal paths are of that form. For example, any path that has an N in it cannot be minimal, because it has to include a superfluous S to cancel the N and arrive at the escape square. Now take any minimal path. Since it includes only S and E moves, the RemoveBackup operation can never be applied to it, since that requires a pair of moves like EW or NS. Flip can apply, but it does not change the moves in the trace, but only changes their order, giving a trace that is also minimal.

While this argument is adequate for our simple cage, it does not cover all situations, for example those in which some barrier within the cage needs to be skirted. Appendix C discusses these more complex situations.

To understand property (2), consider first a nonminimal path that does not include a loop. For ease of exposition, let's imagine quite a small cage, 4 squares by 4, and suppose the starting square is at the northwest corner, and the escape square at the southeast corner. Any path that includes just 4 E moves, and 4 S moves, will lead to escape, and has minimum length. The path SEENEESSSS leads to escape, but it is not minimal, since it includes 10 moves, where 8 suffice. There's an unnecessary N move, and that must be accompanied by an extra S move to cancel its effect. If the extra N were next to the extra S, RemoveBackup would apply, potentially removing the extra S; this possibility renders this path unstable. Now we note that the Flip operation can move the N operation, and/or the S operations, so that the N and S are adjacent. For example, a sequence of Flips, each changing EN to NE, gives

SEENEESSSS -> SENEEESSSS -> SNEEEESSSS

whence RemoveBackup can produce EEEESSSS, which is minimal. Another possible sequence of Flips is

SEENEESSSS -> SEEENESSSS -> SEEEENSSSS, whence RemoveBackup can give SEEEESSS, which is also minimal. That is, if a sequence contains moves that cancel one another, like NS or EW, Flips can rearrange the sequence in such a way that the canceling moves are adjacent, and subject to elimination by RemoveBackup. Thus such sequences are unstable.

The same argument shows that no path in this cage that forms a loop can be stable. Because any loop runs from any point on the loop back to the same point, it has to contain pairs of moves that cancel.

Barriers within the cage complicate this analysis, because, when these are present, a sequence of moves can contain moves that would cancel if adjacent, but cannot be moved together by Flips. This will happen if one of the squares through which the Flipped path would move is a barrier. In the examples given earlier, the corner ES cannot be Flipped to SE if the square to the south of the start of the sequence is a barrier.

In these situations a minimal path *can* contain canceling moves within it, as needed to skirt a barrier. The discussion in Appendix C covers these cases, arguing that all and only the minimal paths are stable under RemoveBackup and Flip, unless barriers in the cage form islands. In such cases there can be non-minimal paths that are stable, because Flips cannot move canceling moves together, even though the path is not minimal

Appendix C also argues that any non-minimal path can be transformed into a minimal path, by RemoveBackup and Flip, unless barriers in the cage form islands. If an island is present, there will be paths around the island that form loops, and are stable. This happens because the barrier within the loop blocks the Flips that would be needed to move together moves that cancel one another.

Before leaving these escapes, let us examine the question of repeated sequences of movements for the PR model, as we did for the noise model. Because the escape paths are much shorter for this model, we can't compare the same segments as for the noise model (40, 20, and 0 moves before the escape on the last trial) but will look at segments that end 25, 10, and 0 moves before escape. As before, we'll look at 30 trials each from 100 model PR cats. For these, the median number of repetitions is identical for all three choices of sequences: 25. So we again find no support for Guthrie's expectation that sequences of movements immediately before escape are better preserved than earlier ones.

How will the PR model cat behave if placed in a maze, a cage with complicated barriers that create multiple dead ends? Unless the walls of the maze create islands, the analysis just given suggests that the model should learn to traverse the maze efficiently. Any path that explores a dead end in a maze will be unstable, and thus eventually pruned. Figure 7 shows a maze whose design is based on that used in Rosenberg et al. (2021), and Figure 8 (a,b) shows model results. The cats escape much more rapidly on their second and later trials than on their first, and continue to improve, though more slowly, over 20 trials.

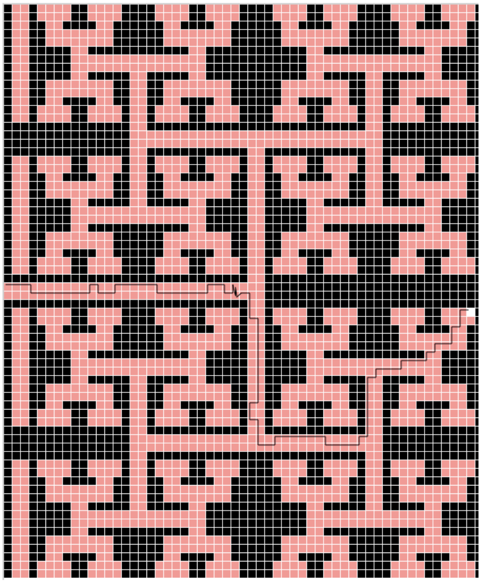

Figure 7: A maze based on that used by Rosenberg et al. (2021). The maze presents a collection of T-junctions, offering an animal a choice of turning right or left. A sequence of seven such choices is needed to move from the maze entrance to the escape square.

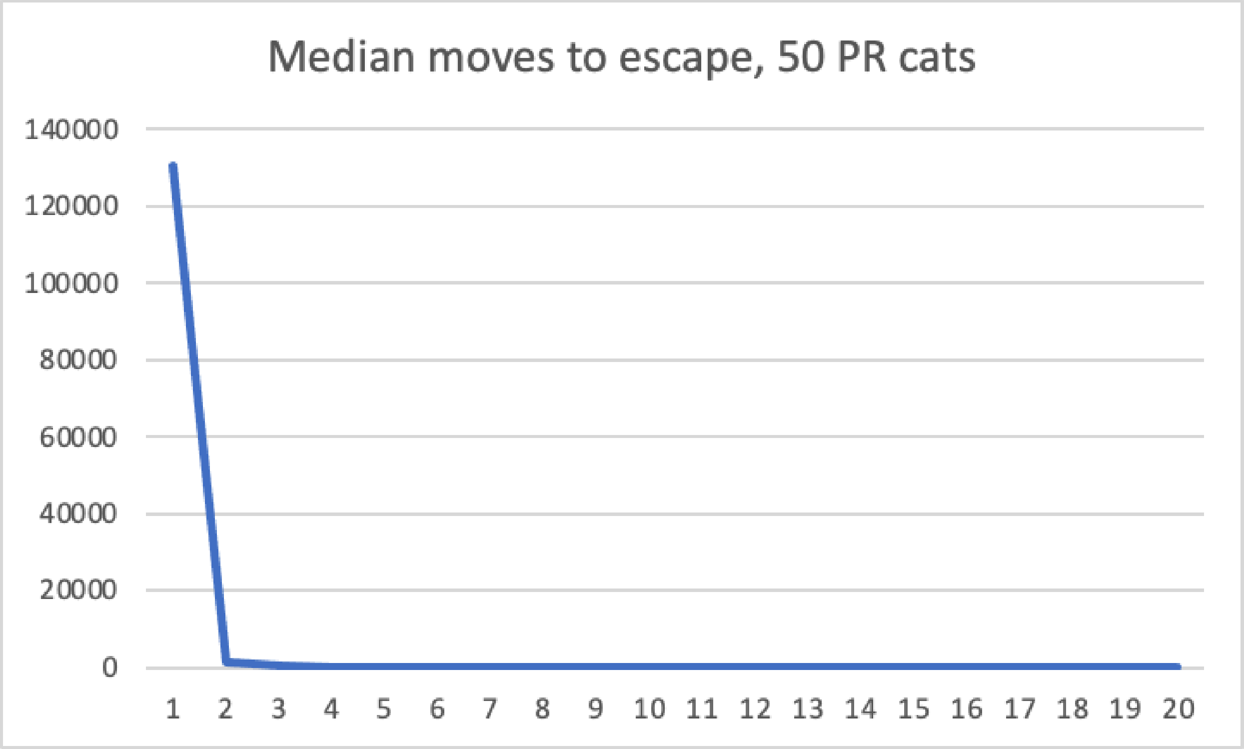

Figure 8a: Median moves to escape from the T-maze shown in Figure 7, for 20 escapes by 50 PR model cats. The median falls from about 130,000 moves for the first escape, to about 1500 on the second escape, to 100 on the eighth escape.

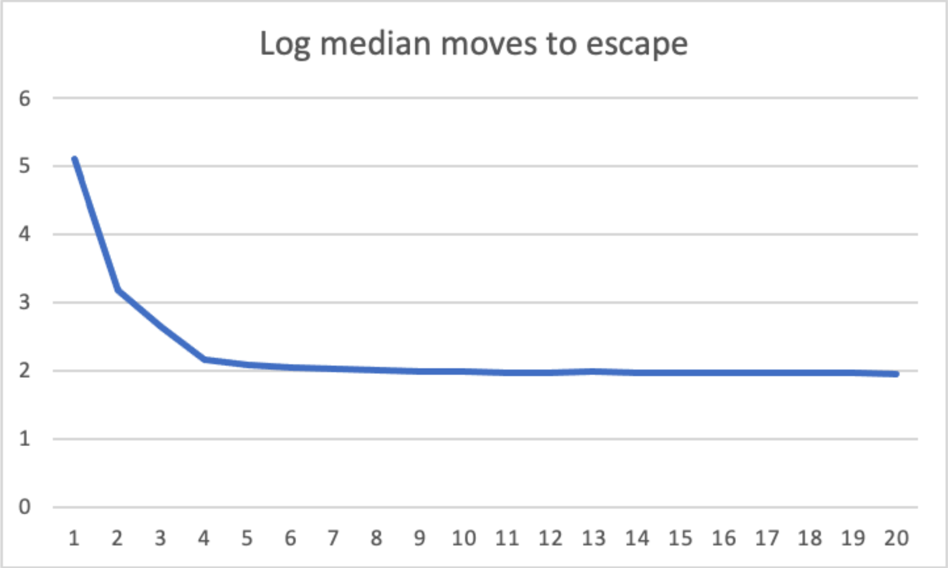

Figure 8b: Logs of the medians shown in Figure 8a. The log scale makes it clearer that the medians are falling quite rapidly over the first four escapes, with much less reduction over the remaining trials.

Eventually the cats attain optimal escape paths. In a simulation of 20 PR cats, over 50 escapes, all twenty have attained a move count of 88, the optimum for this maze, by their 42d escape.

Figure 9 shows a related maze, made so narrow that the cat cannot reverse direction without backing up. In this maze, exiting from a dead end can only happen via RemoveBackup, and this must happen in every case, if the cat is to escape. That is, the maze is learned perfectly in just one escape. Escape data from 50 PR cats in this maze show median moves to escape of about 5700 for the first trial, but, as expected, each escape thereafter, for all 50 cats, takes just 36 moves, the minimum possible for this maze. Appendix D discusses how these results compare to findings reported by Meister and collaborators (Rosenberg et al., 2021), in a study using mice.

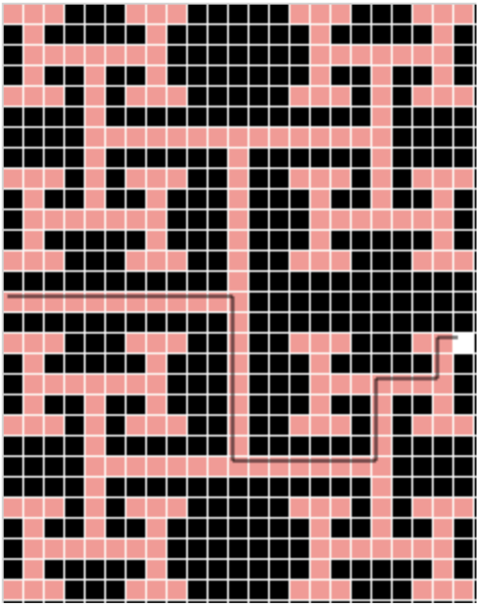

Figure 9: A variant of the T-maze from Rosenberg et al. (2021), with all paths only one square in width.

How does the behavior of the model accord with Guthrie's ideas about maze learning? As already mentioned, Guthrie never focussed on mazes in his research. But he does comment, saying two rather different things.

In Guthrie (1940), he says,

> [W]hatever blinds [dead ends] are entered on the first occasion tend to be entered on the following occasion, and a blind, when it is at last eliminated, tends to be eliminated for the remainder of the series.

Figure 10 (a-e) shows the traces of the first five escapes of a PR model cat in the wider maze shown in Figure 7. As can be seen, there are dead ends entered along the way that are repeated, but eventually eliminated, though they may be reentered.

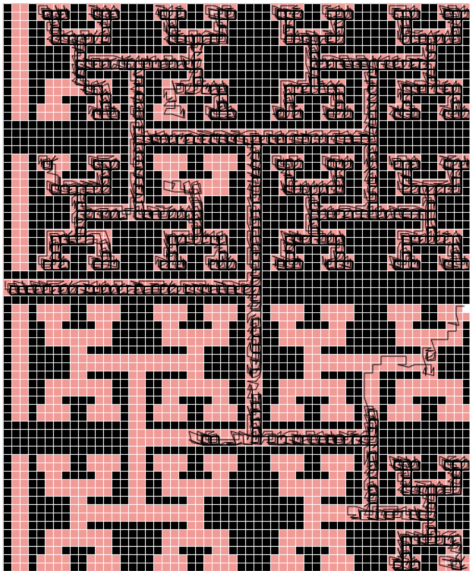

Figure 10(a): Trajectory of first escape of a PR cat in the T-maze of Figure 7. The cat explores the NE quadrant of the maze completely, the NW quadrant largely, and the SE-most sixteenth also completely, while moving only a short way into the SW quadrant. In the process it makes wrong turns at the first, second, third, and fifth T-junctions, before escaping (as well as making various irrelevant moves in parts of the maze that are off the optimal path to escape.)

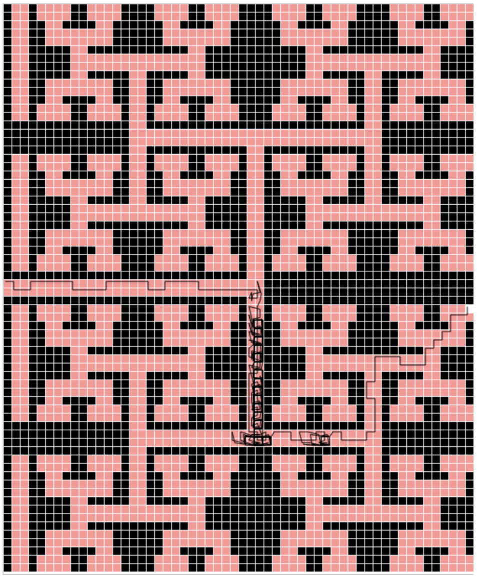

Figure 10(b): Trajectory of second escape. Cat visits much less of the maze, making wrong turns only at T-junctions 1 and 2.

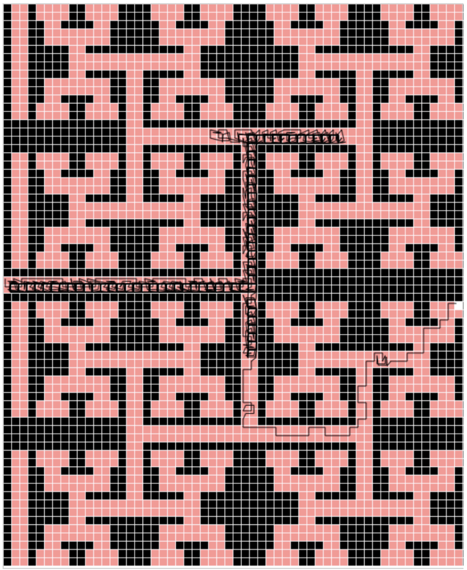

Figure 10(c): Third escape. Wrong turns at T-junction 1, with re-exploration of portions of the maze reached by that turn.

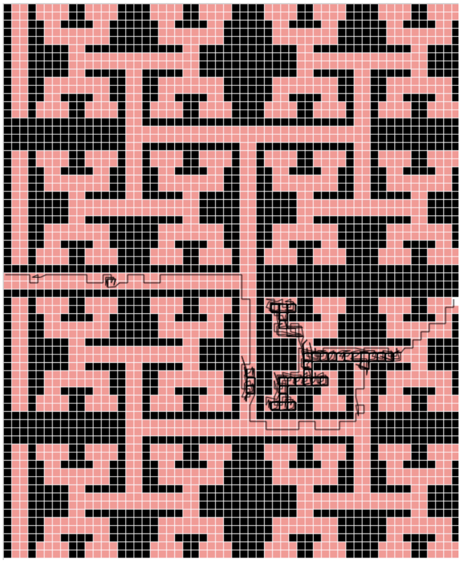

Figure 10(d): Fourth escape. Wrong turn at T-junction 4, with exploration of portions of the maze reached from that turn.

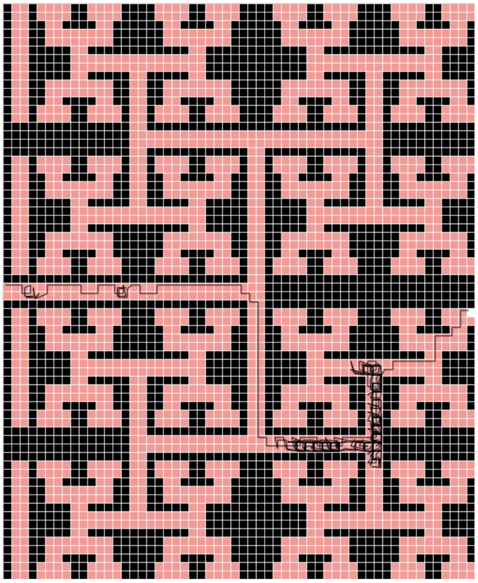

Figure 10(e): Fifth escape. Wrong turns at T-junctions 3 and 4. No long side explorations.

In *The Psychology of Learning* (1935) Guthrie talks about the elimination of dead ends in a different way, suggesting that the fact that the animal leaves the dead end has to lead to its not entering the dead end again:

> Such useless movements as entering a blind alley tend to be eliminated not because that is a general law of learning, but because blind alleys are left as well as entered, and always left after they are entered.

This is the behavior we see in the narrow maze of Figure 9. As discussed above, all dead ends are eliminated in a single escape. Thus we can credit Guthrie with imagining how behavior is different for mazes with passages of different widths, translating his account of leaving a dead end into the action of the RemoveBackup operation in a narrow maze.

Summing up, the PR model shows that escape or maze learning can happen in much the way Guthrie envisioned. Reward plays no role in the learning process, and associations between locations and movements are formed in an ungraded, all or nothing way. In some cases, notably the narrow maze shown in Figure 9, the model learns a minimal escape path on its first success, just as Guthrie thought.

As we have seen, obtaining these results required some ideas that neither Guthrie nor his critics suggested. Some means must be provided to break the loops that occur very frequently when movements are simply stored and repeated. Likely Guthrie would have argued that such loops would be infrequent, in a state space of realistic complexity, in which fatigue and other factors might make an actual repetition of a state unlikely. But then it is unclear what would be learned, if states are never revisited. Perhaps a solution could be found by representing a situation as a bundle of attributes, with associations between these attributes and movements being stored, and some means found for choosing a movement when these associations conflict. This is the direction followed by Estes in his development of Guthrie's ideas, as mentioned earlier, and anticipated in some respects by Guthrie's former student, Virginia Voeks(1948)

The approach developed here has the advantage that it accounts for a feature of escape behavior to which Guthrie deliberately did not attend: the tendency of escape to occur more and more efficiently. Accomplishing this seems to require some representation of the geometry of the state space, so that longer and shorter paths can be related. Perhaps geometric relationships between squares could be reflected in attributes shared between neighboring squares, in an Estes-like approach, but how this sharing could result in path shortening, and removal of loops, would remain to be worked out.

Guthrie's insight seems to have failed him in thinking about the effects of escape from the learning situation. He argued that escape protected recent actions from interference. But we've seen that opportunity for interference seems hardly to be affected by escape, in the context of a sequence of escapes, with the animal being put back in the cage after each escape. Further, the model results suggest that actions near the escape are not more likely to be repeated than earlier actions.

On the other hand, it does seem that there was something right in Guthrie's focus on escape. Suppose we mark some square that the model animal reaches in the cage, but we allow it to continue to run about for a while, before taking it out of the cage, and restarting it. There is no guarantee that the animal will again reach that marked square, but there is a guarantee that it will again reach the square it was in when we took it out of the cage. This happens because of our storing in each square the last move the animal took. What is happening is not that moves near the escape are protected from interference, but that the *whole trace* of moves leading to escape, the last moves taken in each square, is preserved from one escape to the next. Actually, under the action of RemoveBackup and Flip, it is not the trace itself that is preserved, in general, but a modified version of the trace, that still leads to escape.

Stepping back, one can regret that Guthrie and his disputants lacked the means to create computational models with which to explore the potential of their ideas. Virginia Voeks made a brave effort to use formal analysis to sharpen up Guthrie's ideas (Voeks, 1950), but could not represent other than very simple situations.

It isn't very clear, however, what difference the kind of results presented here would have made to Guthrie and his contemporaries. While such results can illuminate the logical potential of a set of ideas, they cannot establish that these ideas really capture what is happening in animals. Of course the disputants performed many observations of animals, varying aspects of situations in ways they felt would give results that would cut for or against different theoretical positions. But as Jensen observes, these efforts, too, on top of the extensive verbal argumentation of the time, settled little.

The work of Meister and collaborators, mentioned earlier, aims ultimately to link maze behavior to neural activity. As such work proceeds, perhaps clarity will emerge. In the process, understanding of the logical potential of ideas, including those of Guthrie, may be helpful.

References

Benvenutti, Radharani, Matheus Marcon, Matheus Gallas-Lopes, Anna Julie de Melo, Ana Paula Herrmann, and Angelo Piato. "Swimming in the maze: An overview of maze apparatuses and protocols to assess zebrafish behavior." *Neuroscience & Biobehavioral Reviews* (2021).

Bower, G. H. "A Turning Point in Mathematical Learning Theory." *Psychological review* 101, no. 2 (1994): 290-300.

Eich, Eric. "Theoretical issues in state dependent memory." *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (1989): 331-354.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review, 57*(2), 94–107.

Garcia, John. "Tilting at the paper mills of academe." *American Psychologist* 36, no. 2 (1981): 149.

Guthrie, Edwin Ray. "Psychology of Learning." New York: Harper (1935).

Guthrie, Edwin R. "The effect of outcome on learning." *Psychological Review* 46, no. 5 (1939): 480.

Guthrie, Edwin R. "Association and the law of effect." *Psychological Review* 47, no. 2 (1940): 127.

Guthrie, Edwin Ray, and George Plant Horton. "Cats in a puzzle box." Rinehart, 1946.

Jensen, Robert. "Behaviorism, latent learning, and cognitive maps: needed revisions in introductory psychology textbooks." The Behavior Analyst 29, no. 2 (2006): 187-209.

Moore, Bruce R., and Susan Stuttard. "Dr. Guthrie and Felis domesticus or: Tripping over the cat." *Science* 205, no. 4410 (1979): 1031-1033.

Rosenberg, Matthew, Tony Zhang, Pietro Perona, and Markus Meister. "Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration." *Elife* 10 (2021): e66175.

Sheffield, Fred. D. "Edwin Ray Guthrie: 1886-1959." The American Journal of Psychology , Dec., 1959, Vol. 72, No. 4 (Dec., 1959), pp. 642-650: 642-650.

Small, Willard S. "Experimental study of the mental processes of the rat. II." *The American Journal of Psychology* (1901): 206-239.

Thomas, Elizabeth, Peter J. Snyder, Robert H. Pietrzak, and Paul Maruff. "Behavior at the choice point: decision making in hidden pathway maze learning." *Neuropsychology review* 24, no. 4 (2014): 514-536.

Thorndike, Edward L. "Animal intelligence: An experimental study of the associative processes in animals." *The Psychological Review: Monograph Supplements* 2, no. 4 (1898): i.

Voeks, Virginia W. "Postremity, recency, and frequency as bases for prediction in the maze situation." *Journal of experimental psychology* 38, no. 5 (1948): 495.

Voeks, Virginia W. "Formalization and clarification of a theory of learning." *The Journal of Psychology* 30, no. 2 (1950): 341-362.

Zhang, Tony, Matthew Rosenberg, Pietro Perona, and Markus Meister. "Endotaxis: A Universal Algorithm for Mapping, Goal-Learning, and Navigation." bioRxiv (2021) https://www.biorxiv.org/content/10.1101/2021.09.24.461751v1

Appendix A. Traces of escapes specify paths that lead to escape, and hence contain no loops

Consider a model cat moving from the starting square, and ending at some other square, S, carrying out a sequence of movements M. The trace of the cat's movement is the record created by storing, in each square the cat visits, the move the cat last made, when in that square. The path specified by a trace is obtained by starting the cat in the starting square, and carrying out the stored moves in each square, as it encounters them. We will prove by induction on the length of M that the trace of M specifies a path that reaches S

For the base case of the induction, let M be of length one. In this case S was reached by a single move from the start square, and that move will have been stored in the start square. So the path specified by the trace will take the cat to S, as required.

The inductive hypothesis is that for all sequences M of length (say) n or less, the path specified by the trace of M takes the cat to the square S reached by M. Now consider a sequence of moves M’ of length n+1, that includes one more move, taking the cat from S to another square, S’. By hypothesis, the trace of M specifies a path that takes the cat to S. The move stored there is the move that takes the cat to S’, so the trace of M’ specifies a trace that takes the cat to S’, as required.

It follows that the trace of any escape cannot contain a loop. For if it did, the path specified would not reach the escape square, as we have seen that it must.

Appendix B. About the movements of cats.

The report by Guthrie and Horton (1946) dealt largely with observations of repeated patterns of movements and poses of escaping cats, illustrated by tracings of frames taken from films of the escapes, capturing the moment a cat triggered a mechanism that unlatched the door of its cage. These tracings showed, indeed, that the pose of a cat in one escape was often very much like its pose in other escapes, as Guthrie and Horton wished to demonstrate.

In a paper in *Science*, in 1979, “Doctor Guthrie and Felis domesticus or: Tripping over the cat,” Bruce Moore and Susan Stuttard argued that Guthrie and Horton had grossly misunderstood the meaning of these observations. In many cases, the escape trigger was a vertical rod, against which the cats would rub, so many of the escape frames show the cats rubbing against the rod. Moore and Stuttart knew that rubbing against vertical objects, in the presence of people, is a common behavior in cats, and asserted that this behavior was not something learned by the cats in the Guthrie and Horton study at all, but was already in the cats’ behavioral repertoire.

This critique, with its catchy, if mocking, title, has had considerable impact. Here is Revulsky (1985):

> The work of E.R. Guthrie (1952), the worst example of a learning theorist dedicated to the proposition that animals function like extremely simple machines, illustrates the same point. The main evidence he offered to support his approach to learning was observations of the behavior of cats (Guthrie and Horton, 1946). Moore and Stuttard (1979) have shown that these cats were not learning at all but exhibiting an instinctive rubbing reaction to the presence of humans. In retrospect it is hard to imagine how anybody with any acquaintance with cats could have made the mistake Guthrie did. Guthrie’s theory is still regarded as being of sufficient importance to merit a chapter in the most highly regarded textbook on learning theory (Bower & Hilgard, 1981).
>
> Only within a tradition in which biological factors were almost totally ignored could errors like these be perpetrated for decades.

Moore and Stuttard were also commended, and Guthrie ridiculed, in a Distinguished Scientific Contributions award address by John Garcia (1981).

This scathing critique, by Garcia, Revulsky, Moore and Stuttard, seems to overlook some aspects of Guthrie and Horton's observations. First, the repeated poses are not simply instances of rubbing, but show very specific manners of rubbing, with one cat's rub being quite different from another's. Second, and more concerning, in many of Guthrie and Horton's trials, the trigger was not a vertical rod, but a photocell beam that could not be rubbed on at all. These escapes, too, show consistent poses within, but not between, cats.

The critics seem to have imagined that Guthrie's aim was to explain rubbing as learned, when it occurred. His actual purpose was to show that individual cats repeated individual details of their behavior, when in repeated situations.

Appendix C. Path reduction operations.

The RemoveBackup and Flip operations act so as to modify a *trace*, that is, the moves stored in the squares of a cage or maze. Their effect is to reduce the length of a path specified by a trace, if reduction is possible.

Recall that the RemoveBackup operation acts on a pair of moves that cancel one another, such as NS. When the model cat moves N into the square where S is stored, the stored S is replaced by a random permissible move. Two special cases are of particular interest.

First, suppose that the square with the stored S is surrounded by barriers, forming a dead end, in which S is the only permissible move. Then the S will not be changed, and the cat will return to the square with the N. This triggers RemoveBackup again. This time the N could be left as it is, so that the cat is moving back and forth, but eventually the N will be replaced by some other move. When this happens the dead end has been avoided.

Second, suppose there is another permissible move from the square with the S in it. If this move is used to replace the S, the loop NSNS... is broken, and the cat moves to some new square.

The Flip operation changes a corner like EN into one like NE, by storing N in the square where the E was found, and storing E in the square north of the square that now contains the N. It can only do that if the new corner square does not contain a barrier; if there is a barrier there, Flip does nothing.

Now consider a sequence of moves that takes the cat from some square A to some square B. We will call a path *minimal* if there is no other path from A to B that takes fewer moves. We will not allow paths that enter barrier squares. If a trace specifies a minimal path, we will call the trace minimal.

We will call a trace *stable* if the action of RemoveBackup and Flip changes only the order of moves in the path specified by a trace, but does not introduce any new moves, or eliminate any moves.

Proposition (1): Minimal traces are stable.

In the absence of barriers between the start and end of a path specified by a trace, any minimal trace includes only one or two kinds of moves, say S and E, if the end is anywhere southeast of the start, or just S, if the end is due south. RemoveBackup can never act on such a trace, as it requires pairs of canceling moves, like SN or WE. Flip only changes the order of moves in a trace, so such a trace is stable.

A barrier can require a minimal trace to include opposing moves, so as to skirt the barrier. But Flip can't move these opposing moves together, because the barrier intervenes. If there were no barrier, the trace could be shortened, and hence was not minimal.

Proposition (2): If a trace is not minimal, in a cage that contains no islands, RemoveBackup and Flip can transform it to a minimal trace.

Suppose there are two traces from start to end, one of which is minimal. Because the traces have the same start and end, the non-minimal trace must include the same moves as the minimal trace, plus some number of pairs of canceling moves. The moves in these pairs can be moved together, and eliminated, unless a barrier intervenes. But such a barrier would lie between the paths specified by the two traces, and hence is part of an island. In the absence of islands, these pairs can be eliminated, transforming one trace into the other.

Proposition (3): If a cage contains no islands, all stable traces are minimal.

If a cage does contain an island, there will in general be stable traces that are not minimal. For example, there could be a trace from start to end that goes north of an island, and cannot be reduced by Flip because of the island, and is hence stable, but also a trace that goes south of the island, that happens to be shorter. That means that the northerly trace, while stable, is not minimal.

If there are no islands, any non-minimal trace cannot be stable, since by Proposition (2) it can be transformed into a trace with different moves.

Together, these propositions imply that the traces of a repeatedly escaping cat will eventually become minimal, in the absence of islands. The trace can never become stable, without becoming minimal. There is a sequence of transformations that will transform it to a stable form, and this sequence will eventually be found.

Appendix D. Meister's mice in mazes.

Markus Meister and collaborators (Rosenberg et al., 2021) have collected extensive data on mice exploring a maze like those in Figures 7 and 9. In their studies, slightly water deprived mice are allowed to explore a maze in which a water dropper is installed at one of the leaves. The dropper delivers a single drop of water, and then waits for about 90 seconds before dispensing another drop. Rosenberg et al. observe that the mice learn to run directly to the dropper, and then return to their cage, outside the maze, before repeating.

They note the following:

(1) After considerable exploration on their first few trips into the maze, after just a few visits mice go directly to the dropper.
(2) Even on their first visit, mice return from the dropper to the maze exit without diversion.

Considering finding (2) first, it seems likely that this first-trial mastery reflects a turning bias. It is a feature of the geometry of the maze that a mouse that always takes a turn when it is offered will reach the maze entrance from any point in the maze, without error. Turning bias is of no help for mice going into the maze, however, because each choice is a T, requiring a choice between two turns. Meister (personal communication, 2020) notes that geometric biases of this kind are not unknown; for example, some ants lay out pheromone trails in such a way that, when a junction is reached, the turn requiring the smaller turn always leads to the nest. Such a turning bias in mice could arise because of the value of maintaining whisker contact with walls, or, parallel to the ant case, because mice are somehow able to construct paths of such form as makes the bias useful for navigation. In discussion, Meister notes that in their data, mice do NOT show a turning bias when not returning to the maze exit. Indeed, their data suggest that mice have three different modes of behavior, seeking water, exploring, and exiting the maze. We discuss the exploration mode further, below.

The mice in Rosenberg et al. are rewarded when they find the dropper, rather than escaping from the maze at that point, so some modification to the PR model is needed to compare to this behavior. This can be done by splitting the state space, so that a square visited by a thirsty mouse, on the way into the maze, is treated as distinct from a square visited by a mouse that is not thirsty, on its way out of the maze. That is, the same square can have two stored moves, a thirsty move and a non-thirsty move. Making the distinction introduces a form of state-based learning (see Eich, 1989), in which things learned in one mental or physiological state are more readily retrieved in that state than in different states.

If this is done, the model results shown in Figures 8 and 10 represent the behavior of mice entering the maze. These are in broad accord with finding (1). As Guthrie suggested should be so, and as the results show, the mice seem to learn to find the dropper in just a few trials.

The return paths do *not* agree with the observation (2) in Rosenberg et al., however, unless the turning bias discussed earlier is introduced, in which case, of course, the mice get out of the maze without error. Without the turning bias the unthirsty mice face a challenging exploration on the way out, as their thirsty selves did on the way in. Meister has pointed out a further possibility: The turning bias could be dispensed with, by having the unthirsty mouse simply follow its stored trace backwards. As Rosenberg at al. report, exiting mice do not simply reverse their paths into the maze, but neither would mice following their own trace backwards. The trace eliminates side branches that have been backed out of, in the narrow maze.

As mentioned earlier, the mice in Rosenberg et al. have a third mode of behavior in the maze, exploration. Indeed, they spend more time in this mode than in seeking the dropper, or exiting the maze. This is a sharp departure from the behavior of the PR model. PR does essentially no exploration, after finding the dropper; it only “explores” local adjustments to the path it has found to the dropper, and does not visit other areas of the maze.

The observations about islands in a cage made in Appendix C are of some interest in connection with Meister’s studies. As can be seen, Meister’s maze has no islands, and hence no stable loops can occur. As of this writing Meister and collaborators intend to study a new maze apparatus in which islands can occur. It will be interesting to see if looping behavior shows up for these mazes, as the PR model predicts.

Less dramatically, islands open the possibility of stable paths that are not minimal, meaning that a mouse might revisit the dropper indefinitely often, via a route that is not as short as possible.

In more recent work (Zhang, et al., 2021) Meister's group proposes an *endotaxis* model, with neural implementation, in which exploration plays a key role in allowing the mouse to actually find optimal paths to the dropper, even in the presence of islands. The same machinery also supports efficient exits. This model manages virtual markers, and information about them, in a way that goes well beyond PR and its Guthrie-inspired ideas.

Acknowledgements.

I thank Markus Meister and Owen Lewis for helpful discussions. My interest in Guthrie's work dates back to my graduate student days at the Human Performance Center, University of Michigan, where John Anderson, Jim Greeno, David Krantz, Gary Olson, Judy Olson, Bob Pachella, Dan Weintraub and others created a most stimulating environment.