

Dynamical trajectories in category learning

SHAWN W. ELL

University of California, Berkeley, California

and

F. GREGORY ASHBY

University of California, Santa Barbara, California

Category learning has traditionally been studied by examining how percentage correct changes with experience (i.e., in the form of learning curves). An alternative and more powerful approach is to examine dynamical learning trajectories—that is, to examine how the parameters that describe the current state of the model change with experience. We describe results from a new experimental paradigm in which empirical-learning trajectories are directly observable. In these experiments, participants learned two categories of spatial position, and they were constrained to identify and use a linear decision bound on every trial. The dependent variables of principal interest were the slope and the intercept of the bound used on each trial. Data from two experiments supported the following conclusions. (1) Gradient descent provided a poor description of the empirical trajectories. (2) The magnitude of changes in decision strategy decreased with experience at a rate that was faster than that predicted by gradient descent. (3) Learning curves suffered from substantial identifiability problems.

Traditionally, category learning has been investigated by examining how response accuracy changes with experience. Often, such data are presented in the form of a learning curve, which plots proportion correct against trial or block number. Learning curves are a good non-parametric method for investigating category learning, because no model needs to be specified during their construction. Learning curves are also relatively simple to compute and often provide an effective method for comparing task difficulty across different conditions of an experiment (e.g., Shepard, Hovland, & Jenkins, 1961).

On the other hand, the use of learning curves to test among competing models is severely limited because, in most cases, a variety of different models will be capable of predicting the same learning curves. For example, a popular assumption of many category-learning models is that the trial-by-trial learning of categories is accomplished via a process of gradient descent on the error surface (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Estes 1993, 1994; Estes, Campbell, Hatsopoulos, & Hurwitz, 1989; Gluck & Bower, 1988; Hurwitz, 1990; Kruschke, 1992; Nosofsky, Kruschke, & McKinley, 1992). Roughly speaking, gra-

dient descent algorithms (e.g., the delta rule, back-propagation) predict that the learner will update the parameters that govern learning on a trial-by-trial basis in a way that causes the probability of error to decrease in the fastest possible manner (e.g., Rumelhart, Hinton, & Williams, 1986). Even so, for any increasing learning curve, it will always be possible to find two models that provide perfect fits, where one assumes gradient descent and the other does not (an example will be shown below).

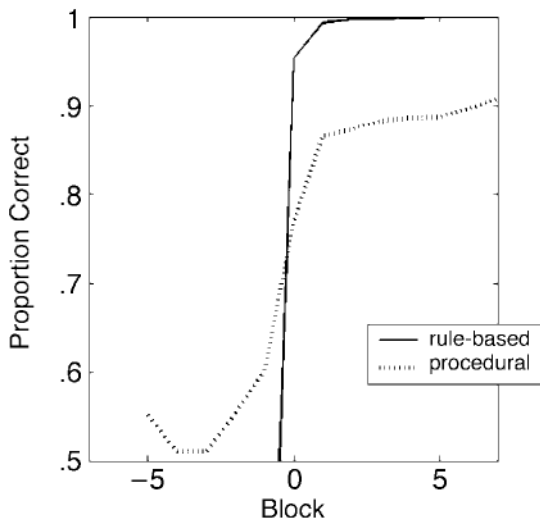
This article has a number of goals. First, we will demonstrate some of the limitations of learning curves when category-learning data are analyzed. Next, we will consider some advantages of considering category learning as a dynamical process that can be studied with classical techniques from dynamical systems theory. In the third section, we will propose a new empirical method in which dynamical trajectories are directly observable. Finally, two experiments will be presented that demonstrate the application of a dynamical systems approach to category learning. Among other results, we will show that the empirical trajectories from our experiments are incompatible with the assumption of gradient descent.

Some Limitations of Learning Curves

Consider an experiment in which observers are asked to classify stimuli into two categories, A and B, and suppose that the stimuli vary across trials on three continuous-valued dimensions (e.g., color, shape, and size). Suppose further that the category structures are such that size is the only relevant dimension. For example, large stimuli might be in Category A, and small stimuli in Category B. In this case, there are a number of ways these categories could be learned. One possibility is that observers use an

This research was supported in part by National Science Foundation Grant BCS99-75037. Portions of this research were reported at the 1999 meetings of the Society for Mathematical Psychology. The authors thank Todd Maddox, Eric-Jan Wagenmakers, Richard Schweickert, In Jae Myung, and an anonymous reviewer for helpful comments on earlier drafts of this article. Correspondence concerning this article should be addressed to S. W. Ell, Psychology Department, University of California, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650 (e-mail: shawnell@socrates.berkeley.edu).

A) Continuous-Valued Dimensions



B) Binary-Valued Dimensions




Figure 1. Simulated backward-learning curves for the rule-based and procedural-based models in category-learning tasks with (A) continuous- and (B) binary-valued dimensions.

explicit rule-based strategy in which they only try unidimensional rules that are easy to verbalize. If we denote the stimulus values on the three dimensions by x_1 , x_2 , and x_3 , then under this scenario, all rules are of the form

Respond A if $x_i > \varepsilon_i$; otherwise respond B,

where ε_i is the decision criterion on dimension i . To use this strategy successfully, an observer must find the relevant dimension and determine the optimal value of ε_i . The latter problem is often referred to as *critical learning*.

A very different possibility is that observers use a strategy in which they integrate information from all the stimulus dimensions at some predecisional stage (e.g.,

by using a form of procedural learning). In this type of learning, each stimulus dimension is assigned a weight, and the category response is determined by computing a weighted combination of the values on each stimulus dimension. For example, an observer using such a strategy might respond according to the following decision rule:

Respond A if $\alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_3 > \beta$;

otherwise respond B,

where α_i is the weight on dimension i and β is the decision criterion. To learn the optimal form of this linear discriminant function, the dimension weights and the decision criterion are optimized in a trial-by-trial incremental process.

The rule-based strategy will produce learning that is nearly all-or-none. If the observer has selected the incorrect stimulus dimension, accuracy will be near chance. On the first trial in which the correct dimension is selected, accuracy will jump dramatically, and the only thing preventing optimal responding will be a suboptimal decision criterion. In contrast, we expect the procedural-learning process to lead to incremental improvements in accuracy—small changes in the dimensional weights should lead to small changes in accuracy. As such, accuracy should gradually rise from chance to near-optimal levels. These two strategies, therefore, should produce the most discrepant possible learning curves.

To verify this intuition, we conducted simulations of a model in which a rule-based decision strategy was assumed and one in which a procedural-learning strategy was assumed.¹ Average simulated backward-learning curves are plotted in Figure 1A.² As was expected, the two models predict quite different backward-learning curves. The rule-based model (solid line) initially performs at chance and then jumps to near-perfect accuracy during the block in which a criterion of 10 consecutive correct responses was met (arbitrarily numbered Block 0). On the other hand, the procedural-learning model (dotted line) demonstrates incremental learning. The initial performance of the procedural-learning model is at chance, but instead of a large jump in accuracy from one block to the next, accuracy gradually improves across a number of blocks.

The simulations shown in Figure 1A assumed continuous-valued stimulus dimensions. What would happen instead if the stimuli from the two categories varied across trials on three binary-valued dimensions? In this case, the stimulus coordinates would form a cube, and each of the two categories would contain four exemplars whose coordinates form opposite faces of the cube. The results of this simulation are shown in Figure 1B. Note that the rule-based model (solid line) and the procedural-learning model (dotted line) initially perform at chance, but then the performance of both dramatically jumps to near-perfect accuracy within a few blocks.

In this experiment, it would be virtually impossible to test between all-or-none and incremental learning by examining learning curves. The problem is that, although the procedural-learning model incrementally adjusts the position of its decision bound (i.e., a plane in this case),

most of these changes have no effect on accuracy, because of the sparse nature of the stimulus space. For example, suppose that on trial n , the decision plane of the procedural-learning model is such that both sides contain two exemplars from Category A and two from Category B. In this case, accuracy will be at chance. Now suppose that the plane is incrementally adjusted so that it moves in the direction of the optimal decision bound. Unless this movement passes through the coordinates of one of the four incorrectly classified exemplars (and does not pass through the coordinates of one of the four exemplars that were correctly classified), accuracy on trial $n + 1$ will be the same as that on trial n . Thus, with binary-valued stimulus dimensions, the incremental procedural-learning model predicts that accuracy will remain at chance as the decision bound is incrementally adjusted until the bound rotates through the coordinates of an incorrectly classified exemplar. At this point, accuracy will increase discretely. Figure 1B shows that this identifiability problem is so severe that the incremental procedural-learning model predicts essentially the same learning curves as the all-or-none rule-based model. In other words, simply changing the nature of the dimensions (from continuous to binary valued) produced a lack of identifiability at the level of analysis typically used in category-learning experiments.

Dynamical Trajectories

Every model can be characterized by a set of parameters. Excluding any stochastic influences, specifying the numerical values of these parameters fixes the predictions of the model. Let $\underline{\omega}$ denote the vector of parameters for some particular model. The so-called parameter space of the model contains a dimension for each parameter in the $\underline{\omega}$ vector. For any numerical value of $\underline{\omega}$ —that is, for any point in the parameter space—the model will make specific numerical predictions that will be associated with some overall probability of an incorrect categorization response, $P(\text{error} | \underline{\omega})$. Adding this extra $P(\text{error} | \underline{\omega})$ dimension to the parameter space creates a probability-of-error surface. In any experiment, some particular combination of parameters $\underline{\omega}^*$ will predict fewer errors than any other possible combination will. Thus, the lowest point on the probability-of-error surface will have coordinates $\underline{\omega}^*$, and the elevation there will be $P(\text{error} | \underline{\omega}^*)$. The behavior of the model on the first trial of the experiment can be described by the initial parameters $\underline{\omega}_0$, so learning can be described as movement through the parameter space in an attempt to find the lowest point on the probability-of-error surface—that is, as an attempt to move from the initial point $\underline{\omega}_0$ to the optimal point $\underline{\omega}^*$. The path taken by the model through parameter space is called its dynamical trajectory. A major goal of this article is to show that dynamical trajectories provide a powerful alternative to learning curves as a tool with which to compare competing learning models.

Dynamical trajectories through the parameter space completely characterize the learning assumptions of a

model. In particular, any learning model must define some algorithm that describes how $\underline{\omega}$ changes on a trial-by-trial basis—that is, it must define the allowable dynamical trajectories. In category learning, the most popular approach has been to assume that learning follows a process of gradient descent—in other words, that the trajectories always move in the direction of the steepest possible descent down the probability-of-error surface. The gradient in the probability-of-error surface in every direction is specified in the vector of partial derivatives $\partial P(\text{error} | \underline{\omega}) / \partial \underline{\omega}$. Thus, according to gradient descent algorithms, the change in $\underline{\omega}$ across trials is completely determined by $\partial P(\text{error} | \underline{\omega}) / \partial \underline{\omega}$:

$$d\underline{\omega}/dt = -\alpha \partial P(\text{error} | \underline{\omega}) / \partial \underline{\omega}. \quad (1)$$

In particular, gradient descent algorithms assume that α is a constant. In this model, $\underline{\omega}$ is changed in the direction that causes the greatest possible decrease in $P(\text{error} | \underline{\omega})$, and the steeper the gradient in the error surface at the point $\underline{\omega}$ on trial n (or at time t), the greater the resulting change in $\underline{\omega}$ on trial $n + 1$.

Admittedly, it is unlikely that people perform the calculus required by Equation 1 on a trial-by-trial basis; however, it is reasonable to assume that observers may adjust the parameters that govern learning with the goal of minimizing error. In fact, psychologists consider the assumption of gradient descent to be so reasonable that it is almost a universal feature of current category-learning models (Ashby et al., 1998; Erickson & Kruschke, 1998; Estes, 1993, 1994; Estes et al., 1989; Gluck & Bower, 1988; Hurwitz, 1990; Kruschke, 1992; Nosofsky et al., 1992). Furthermore, if the goal is to accurately describe how people develop category representations, the assumptions that define the manner in which learning occurs (e.g., gradient descent) are as critical as assumptions regarding the storage of information, response selection, and category access.

One especially popular method of implementing gradient descent is to cast the model in the form of a connectionist network. Briefly, these models assume a multilayer network of interconnected nodes. The parameters of the model include the strengths (or weights) of these connections. Information is passed from an input layer to an output layer, sometimes through an intermediate (hidden) layer or layers. The input layer represents the stimulus, and each node in the output layer is associated with a response. Learning in these networks is the process of modifying the strengths of the connections between the layers—almost always via an algorithm that implements gradient descent. Well-known examples include the delta rule and back-propagation, both of which implement gradient descent, but in different types of connectionist networks (Haykin, 1994). Models of this type have been quite successful in accounting for various aspects of human category-learning performance, but the gradient descent assumption has never been directly tested, so it is not clear whether gradient descent accurately describes the nature of learning in human observers.

As learning progresses from the initial state $\underline{\omega}_0$ to the optimal state $\underline{\omega}^*$, changes in $P(\text{error} | \underline{\omega})$ tend to decrease, because in most cases, the probability-of-error surface flattens out around $\underline{\omega}^*$. Gradient descent predicts that the magnitude of changes in $\underline{\omega}$ are proportional to the gradient—that is, gradient descent algorithms make large changes in $\underline{\omega}$ when $\underline{\omega}$ is in a region where the probability-of-error surface is steep and small changes in $\underline{\omega}$ when $\underline{\omega}$ is in a region where the probability-of-error surface is shallow. For this reason, gradient descent typically predicts that the magnitude of changes in $\underline{\omega}$ will tend to decrease as the observer gains experience in the task. At the extreme, these models predict no further changes in $\underline{\omega}$ once $\underline{\omega} = \underline{\omega}^*$, because the probability-of-error surface is flat at $\underline{\omega} = \underline{\omega}^*$ (since $\underline{\omega}^*$ contains the coordinates of the minimum of the probability-of-error surface). In fact, empirical evidence for such decreases has been reported (Busemeyer, Swenson, & Lazarte, 1986). Even so, many category-learning models that incorporate gradient descent include an additional assumption that this decrease is even more severe than predicted by standard gradient descent. This assumption, which is known in the literature as *cooling*, is formally implemented by allowing the constant α in Equation 1 to decrease with experience. To our knowledge, the assumption of cooling has never been empirically tested.

Dynamical trajectories provide more identifiability than learning curves do. To see this, consider the trajectories of two different category-learning models that are characterized by the same pair of parameters ω_1 and ω_2 . For example, both models might assume that observers

use a linear decision bound with slope ω_1 and intercept ω_2 in order to categorize stimuli that vary across trials on two dimensions. Note that the rule-based and the procedural-learning models described in the previous section satisfy this description. However, suppose that the two models postulate different learning mechanisms. In particular, suppose that the first model assumes that learning obeys gradient descent, whereas the second postulates a learning process that dramatically violates gradient descent.

Figure 2 shows dynamical trajectories from these two models across a series of six trials in a hypothetical category-learning experiment. The graph is a contour plot of the probability-of-error surface over the (ω_1, ω_2) parameter space. Darker areas in Figure 2 indicate a smaller probability of error, and all the points on the same contour have the same probability of error. The solid white lines are the trajectories of the two models, and the filled circles represent the position of each model on the error surface across trials. Initially, both models begin at approximately the same point on the error surface (i.e., trial 1) and end at the surface minimum, $\underline{\omega}^*$. As is predicted by gradient descent, the trajectory of Model 1 proceeds directly downhill, whereas the trajectory of Model 2 clearly violates gradient descent. Despite this large qualitative difference in the learning assumptions, note that both models predict the same probability of error on each of the six trials illustrated in Figure 2 (e.g., the trajectories of both models on trial 3 lie on the same contour line). As a result, these two very different model trajectories predict identical learning curves. Thus, although it would be easy to test between the models if dy-




Figure 2. State space defined by the parameters ω_1 and ω_2 overlaid upon a contour plot of a probability-of-error surface. The solid white lines represent the model trajectories, and the filled white circles represent each model's position on the error surface across six trials. See the text for details.

namical trajectories were available, they are nonidentifiable with traditional methods of data analysis (i.e., learning curves).

The advantage of learning curves is that they are non-parametric—in the sense that they can be estimated without making any assumptions about the observer's decision strategies. The disadvantage is the high degree of nonidentifiability that exists at this level of data analysis. As we have seen, models that make strikingly different assumptions about decision strategy can predict identical learning curves. In contrast, the identifiability problems are virtually eliminated with a dynamical trajectories analysis. Two models postulating different learning algorithms almost necessarily must predict different dynamical trajectories. A major disadvantage with a dynamical trajectories approach, however, is that it is parametric—that is, since the trajectories are computed through some specific parameter space, trajectories can be computed only if a decision model is first specified. To make matters worse, the numerical values of the parameters that collectively define the state of the model on each trial are unobservable in typical applications. As such, these parameters must be estimated from the data. This is a serious impediment, because during learning the values of these parameters will presumably change on a trial-by-trial basis. One solution to the parameter estimation problem is somehow to make the trajectories directly observable. If this could be done, the parameters would always be known, and the problem of estimating the trajectories would vanish, because they would be empirically observable. In the remainder of this article, we will propose an experimental method that accomplishes this goal, and we will report the results of an experiment in which this method was used.

Making Dynamical Trajectories Directly Observable

There have been several attempts to estimate dynamical learning trajectories at discrete time points during the course of multidimensional perceptual categorization (Alfonso-Reese, 1996; Minda & Smith, 2001; Smith & Minda, 1998, 2000). In none of these studies were the empirical trajectories directly observable. Smith and Minda (1998, 2000; Minda & Smith, 2001) estimated dynamical trajectories predicted by prototype and exemplar models of categorization by fitting these models to relatively small segments of empirical data (e.g., groups of 10 trials; Minda & Smith, 2001). In Alfonso-Reese's experiment, a novel approach, called the *decision bound probe technique* (DBPT), was used to estimate the trajectories. The DBPT interrupts a sequence of traditional category-learning trials (that include trial-by-trial feedback) with probe blocks that occur at various times throughout learning. Each probe block includes 50 trials without feedback, during which participants are instructed to categorize a set of transfer stimuli (that are uniformly distributed across the stimulus space) by using the same decision rule that they used during the

preceding learning phase. The data from the probe blocks is then used to estimate a best-fitting decision bound, whose parameters are assumed to locate the dynamical trajectory within the parameter space on the last learning trial preceding the probe block. Thus, to the extent that participants are able to maintain their decision strategy throughout the probe blocks, the DBPT allows the learning state to be estimated at discrete time points throughout the learning process (i.e., that depend on how many learning trials lie between the probe blocks). Therefore, in the DBPT, as in the work of Smith and Minda, the decision process is observed indirectly, and the trajectories are estimated by making assumptions about the true model of categorization.

In contrast, in several studies, attempts have been made to make the decision process directly observable in categorization tasks in which the stimuli varied across trials on only a single physical dimension. In all these studies, a signal-detection-like decision process was assumed in which the observer sets a criterion on the stimulus dimension and assigns all the stimuli with a value greater than that criterion to one category and all the stimuli with a value less than the criterion to the contrasting category. In this case, the parameter space is one-dimensional (i.e., the value of the decision criterion). A series of studies have developed empirical methods for making the value of this parameter directly observable.

The most common such method is based on the so-called *numerical decision task* (Lee & Janke, 1964, 1965). Briefly, in this task, two univariate normal distributions of numbers are defined, and on each trial, a random sample, x , is drawn from one of the distributions and presented to the observer. The observer's task is to decide whether this number was sampled from Distribution A or B. In most applications, the decision criterion in this task is not known to the participants, but in the *cutoff report* condition in Kubovy and Healy (1977), the criterion, ϵ , was specified by the participants and, therefore, was directly observable. Specifically, on each trial, the participants were first asked to specify a numerical value for ϵ . Next, a random sample, x , from one of the two category distributions was presented, and the participants were to make their responses on the basis of the following decision rule: if $x > \epsilon$, respond A; otherwise, respond B. For example, if $x = 10$ and the participant had set $\epsilon = 5$, the required category response was A. The participants were given feedback about the accuracy of each categorization response, and they were permitted to adjust the value of ϵ on the basis of this information at the beginning of the next trial. The version of the numerical decision task used by Busemeyer and Myung (1992) (Experiment 2) framed the task as a medical classification problem. Here, the observer's choice of ϵ determined his or her diagnosis about the presence or absence of a disease.

Both of these studies were successful in making the decision criterion directly observable on a trial-by-trial basis, although the primary focus was not to investigate

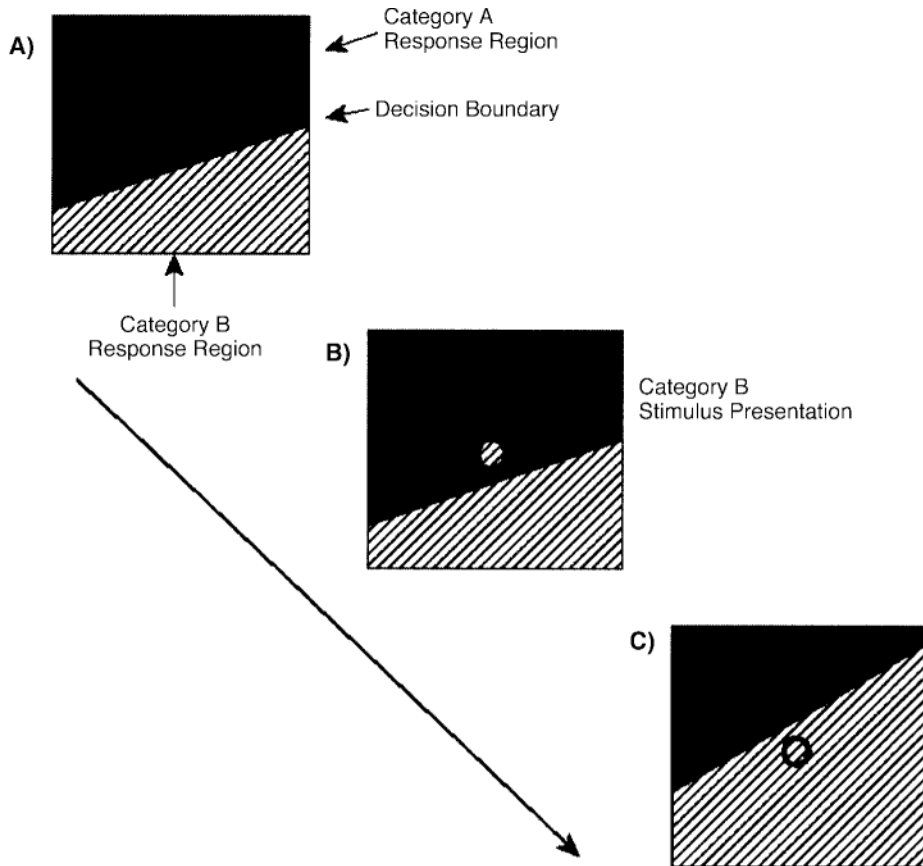


Figure 3. An example of a trial from Experiment 1. (A) At the beginning of a trial, only the response regions are visible. (B) Next, the stimulus is presented, and (C) the observer adjusts the position of the decision bound as needed. Category A and B response regions are represented as solid and hatched regions, respectively.

dynamical trajectories but, rather, to test among various models of criterial learning. Kubovy and Healy (1977) tested between deterministic and probabilistic criterial-learning models, whereas Bussemeyer and Myung (1992) tested the parameter search process of their *hill-climbing* model. Thus, the primary goal of these studies was to test specific models. Although the decision criterion was directly observable, dynamical trajectories were not examined, and the decision task was unidimensional. In contrast, in several studies, dynamical trajectories were estimated in a category-learning task with multidimensional stimuli, but the decision bound was not directly observable on a trial-by-trial basis (Alfonso-Reese, 1996; Minda & Smith, 2001; Smith & Minda, 1998, 2000). The goal of the present experiment was to combine both experimental procedures, so that the decision strategy in a task with multidimensional stimuli was observable on a trial-by-trial basis and this information, in turn, was used to specify dynamical learning trajectories. This experimental technique was used to test the assumptions of gradient descent and cooling and to highlight the nonidentifiability issues associated with learning curves.

EXPERIMENT 1

In our experiment, each stimulus was a single circle that varied across trials in its location on the computer screen. Thus, the stimuli varied on two physical dimensions—vertical and horizontal spatial positions. Exemplars from Category A were characterized by positions on one part of the screen, and exemplars from Category B tended to fall in another part (although the category distributions overlapped). Throughout the experiment, the screen was divided into dark green and dark blue response regions that were separated by a line. An example is shown in Figure 3. The observers could move this line between trials to any position on the screen by pressing appropriate computer keys. Stimulus circles from Category A were displayed in light green, and circles from Category B were displayed in light blue. The observers were instructed to position the line between trials to maximize the probability that light green Category A exemplars fell in the dark green Category A response region and that light blue Category B exemplars fell in the dark blue Category B response region.

The dependent variables of interest were the slope and the intercept of the line that divided the monitor into the two regions.³ For each observer, a plot of these values across trials defined that observer's dynamical learning trajectory for this particular categorization task.

An example showing the sequence of events on a hypothetical trial in which the observer responds incorrectly is shown in Figure 3. The trial begins with only the response regions and the decision bound visible (Figure 3A). Then a Category B exemplar is presented in the Category A response region, thereby indicating an incorrect response (Figure 3B). The participant then makes a response by adjusting the position of the decision bound (Figure 3C) to account for the current stimulus and all other, previously viewed stimuli.

Method

Participants

The observers in all the experiments were 20 undergraduates at the University of California, Santa Barbara, who received partial course credit in an introductory psychology course for participating in the experiment. All the observers reported 20/20 vision or vision corrected to 20/20, reported normal color vision, and completed one session approximately 45 min in duration.

Stimuli and Stimulus Generation

The stimuli in all the experiments were light green (Category A) and light blue (Category B) circles defined by their spatial positions on the monitor. The category structures used in Experiment 1 are shown in Figure 4, along with the initial and the optimal linear decision bounds. The Category A and B stimuli are plotted as plus symbols and circles, respectively. The abscissa represents horizontal position on the monitor, while the ordinate represents vertical position.

Each participant learned four different category structures. The stimuli from all the category structures were generated by randomly sampling from two bivariate normal distributions specified by a mean and variance on each dimension and a covariance between dimensions (see Table 1 for the exact category parameters). The categories always had different means but the same variances and covariance. Under these conditions, the optimal decision bounds are the solid lines plotted in Figure 4. The category structures were chosen so that optimal accuracy was 80%. The order of presentation of the category structures was counterbalanced using a 4×4 Latin square, and each participant was presented with a different pseudo-random ordering of the stimuli in each category structure.

The design of this experiment was based upon the randomization technique (Ashby & Gott, 1988). Briefly, using the parameters listed in Table 1 for Category Structure 1 (CS1), 100 random samples were drawn from both the Category A and the Category B populations. The two samples were then combined to form CS1. The stimuli in the other three category structures were generated by ro-

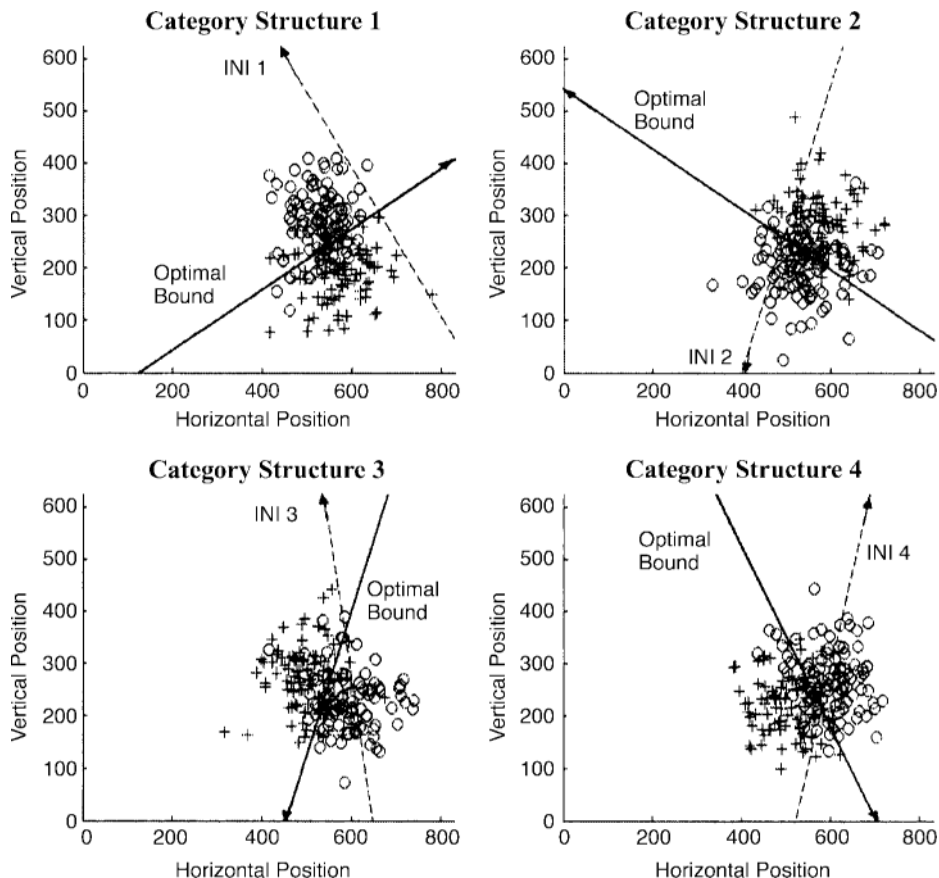


Figure 4. Stimulus distributions for the four category structures used in Experiments 1 and 2. Category A and B exemplars are plotted as plus symbols and circles, respectively. The initial (dashed line) and optimal (solid line) decision bounds are also plotted.

Table 1
Category Parameters of the Four Category Structures Used in Experiments 1 and 2

Category Structure	Means				Variances		Covariance
	Horizontal Position		Vertical Position		Horizontal Position	Vertical Position	
	A	B	A	B			
1	57	525	202	288	3,000	3,000	110
2	575	525	288	202	3,000	3,000	-50
3	503	597	262	228	3,000	3,000	50
4	503	597	228	262	3,000	3,000	-110

tating the CS1 stimulus set: 120° for CS2, 220° for CS3, and 270° for CS4. The parameters for each of these category structures are also given in Table 1.

The stimuli were computer generated and displayed on a View-Sonic 15-in. CRT with 832 × 624 pixel resolution in a dimly lit room. Each random sample (x₁, x₂) from a given category structure defined a pixel point on the monitor and the center of each circular stimulus. All the stimuli subtended a visual angle of approximately 1°. The stimuli were generated and presented using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) in the MATLAB environment.

Four initial decision bounds were chosen (arbitrarily) and were factorially combined with all of the category structures. The initial decision bounds (INI1-4) are plotted with respect to the four category structures in Figure 4. For illustrative purposes, only one initial decision bound was plotted with respect to each category structure.

Display

As was stated earlier, this task was designed so that the decision bound was directly observable. This was accomplished by dividing the monitor into dark green (Category A) and dark blue (Category B) response regions, as depicted in Figure 3 by the solid and the hatched regions, respectively. The decision bound separated these two response regions.

Procedure

The participants were run individually in a dimly lit testing room. On a trial, a stimulus from either Category A (light green circle) or B (light blue circle) was presented in one of the two category response regions. The participant’s task was to adjust the decision bound so that all of the stimuli from a given category structure were divided into two categories as accurately as possible. Specifically, they were instructed to position the decision bound so that as many of the *light green* (Category A) exemplars as possible fell in the *dark green* region (Category A response region) and as many of the *light blue* (Category B) exemplars as possible fell in the *dark blue* response region (Category B). Note that the specific colors used in Experiment 1 were chosen so that the Category A and B stimuli were clearly visible when presented in the Category A and B response regions. The participants were able to move the boundary in six possible ways on every trial—(1) up, (2) down, (3) left, (4) right, (5) rotate clockwise, and (6) rotate counterclockwise—by pressing keys labeled as such on the computer keyboard. The position of the bound (i.e., the slope and the intercept) was recorded after every trial. The participants were told that there were two equally likely categories and were informed of the optimal accuracy. More specifically, the participants were told that the categories would overlap and that it was possible to be incorrect even if the decision bound had been placed in the most accurate position possible.

Visual feedback regarding the current position of the decision bound was given on a trial-by-trial basis. The participants were informed that a correct trial was defined as one in which the stimulus (the center of the stimulus) from Category A (or B) appeared in the Category A (or B) response region. Thus, the stimulus presented on trial *n* provided feedback regarding the position of the bound following trial *n* - 1. When satisfied with the position of the decision

bound, the participants were able to initiate the next trial by pressing the space bar. The current stimulus remained on the monitor until the space bar was pressed. All the trials were separated by a response-stimulus interval of 500 msec. The sequence of events on a hypothetical error trial is shown in Figure 3.

Before beginning the experiment, the participants were given 15 practice trials to ensure that they understood the task. The participants were then told that they would perform the task with four sets of categories and that the stimuli and the optimal decision bound in each of the category structures would be different. The participants were able to move on to the next category structure once they had met the stopping criterion of 80% correct over the last 40 trials in the current category structure. Beginning on Trial 40, this criterion was computed on every trial within a given category structure. An upper limit of 200 trials was imposed for those participants who did not reach the stopping criterion in a category structure. The break between category structures was self-paced.

Results and Discussion

Trials to Criterion

Each of the 20 participants attempted to learn four category structures. Three category structures (out of 80 total) in which the participants did not meet the stopping criterion (80% correct over the last 40 trials) were not considered in the following analyses. The median number of trials to reach the stopping criterion for each category structure is given in Table 2. None of these differences were statistically significant [$\chi^2(3) = 0.806$].

Empirical Trajectories

Recall that the primary dependent variables were the slope and the intercept of the decision bound on every trial and that the decision bound separated the Category A and B response regions. To aid in the visualization of subsequent analyses, the decision bounds were reparameterized as the slope (i.e., degrees of rotation of the bound from horizontal) and the area of the computer screen that was associated with the Category A response region (hereafter referred to as the *area*). In the experiment, the slope ranged from 0° to 360°, where 0° referred

Table 2
Median Trials to Criterion in the Four Category Structures of Experiment 1

Category Structure	Median	Interquartile Range
1	62	57.75
2	82	49.25
3	76	35.75
4	69	45.00

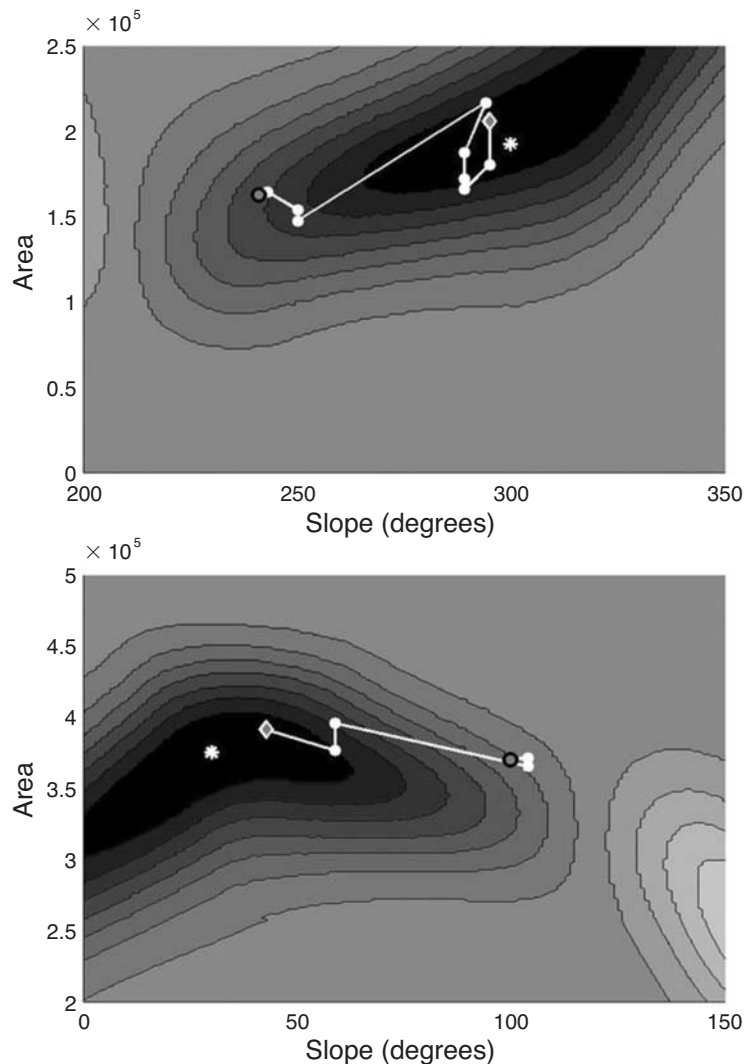


Figure 5. Two representative examples of trajectories (in white) that were classified as not violating gradient descent with respect to the theoretical error surface. The gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound. The top panel is an example from CS4, whereas the bottom panel is an example from CS1.

to a horizontal decision bound in which the Category A response region was above the Category B response region and the slope increased as the bound was rotated in a counterclockwise direction. The area ranged from 0 pixels² to 519,168 pixels² (based on a monitor resolution of 832×624 pixels).

One of the objectives of this experiment was to develop a paradigm in which the dynamical learning trajectories were directly observable. The slope and the area were directly observable on a trial-by-trial basis, so the dynamical trajectories describing the evolution of an observer's decision strategy across trials were directly observable as well. In the plots of these trajectories, pre-

sented below, we overlay the trajectories with contour plots showing the associated probability of error (as in Figure 2; see Ashby, 1992b, for details on computing the error probabilities). It is important to point out, however, that even a model with perfect memory for all category exemplars would not have access to this probability-of-error surface, because its construction requires knowledge of the exact category parameters (i.e., means and covariance matrices). At best, such a model can hope only to estimate these parameters by using some trial-by-trial updating procedure. Therefore, the probability-of-error surface is stochastic across trials, although trial-by-trial changes in that surface should gradually decrease

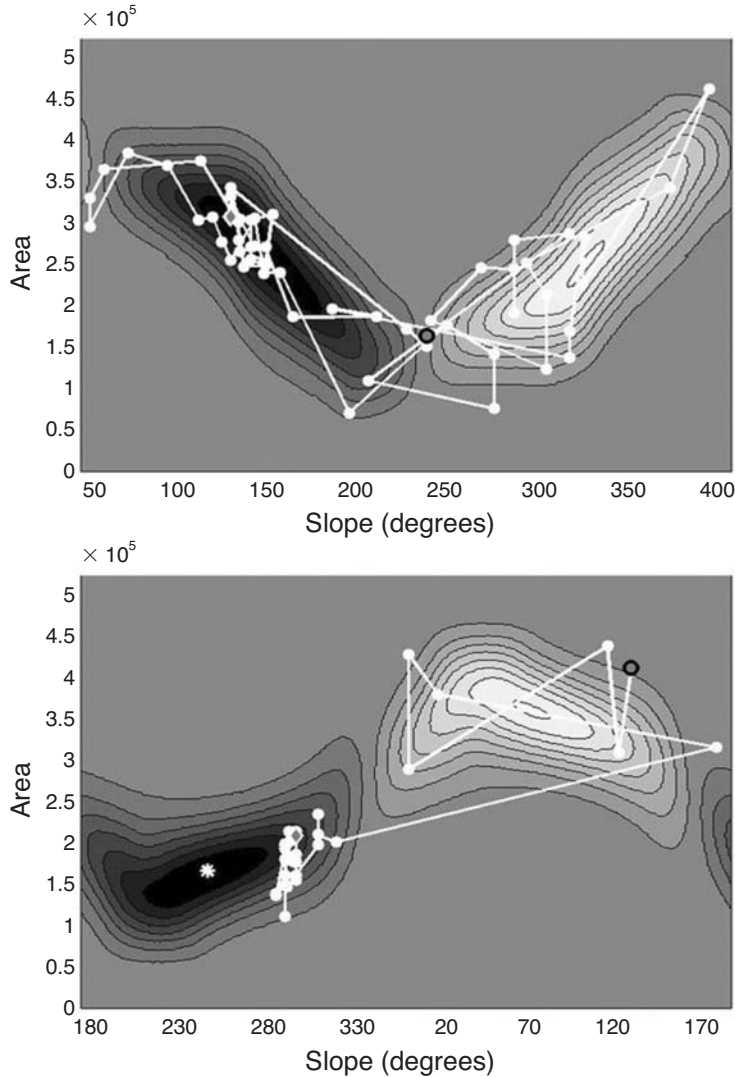


Figure 6. Two representative examples of trajectories (in white) that were classified as violating gradient descent with respect to the theoretical error surface. The gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound. The top panel in an example from CS2, whereas the bottom panel is an example from CS3. The optimal bound is obscured by the trajectory in the top panel and is located at the coordinates 150° , 2.51×10^5 pixels².

with experience. We will examine this issue in more detail shortly.

The Assumption of Gradient Descent

Visual inspection of the empirical trajectories revealed some that appeared to approximately satisfy gradient descent and some that showed apparent dramatic violations. Figure 5 shows two trajectories that are representative of those that were roughly consistent with gradient descent.⁴ Note that in both cases, the trajectories proceed from the initial bound toward the optimal bound in the general direction of steepest descent. On the other hand, Figure 6 demonstrates two trajectories

representative of those that were classified as violating gradient descent—that is, the change in the decision bound is not consistently in the direction of steepest descent. In fact, the change in the decision bound often resulted in an increase, rather than a decrease, in the probability of error. Qualitatively, only a small percentage (approximately 10%) of the trajectories appeared to approximate a gradient-descent-like progression.

Optimal learner. To test the assumption of gradient descent more rigorously, we compared the empirical trajectories with those of an optimal learner model. To construct this model, we added a gradient descent learning algorithm to the general linear classifier (GLC) of deci-

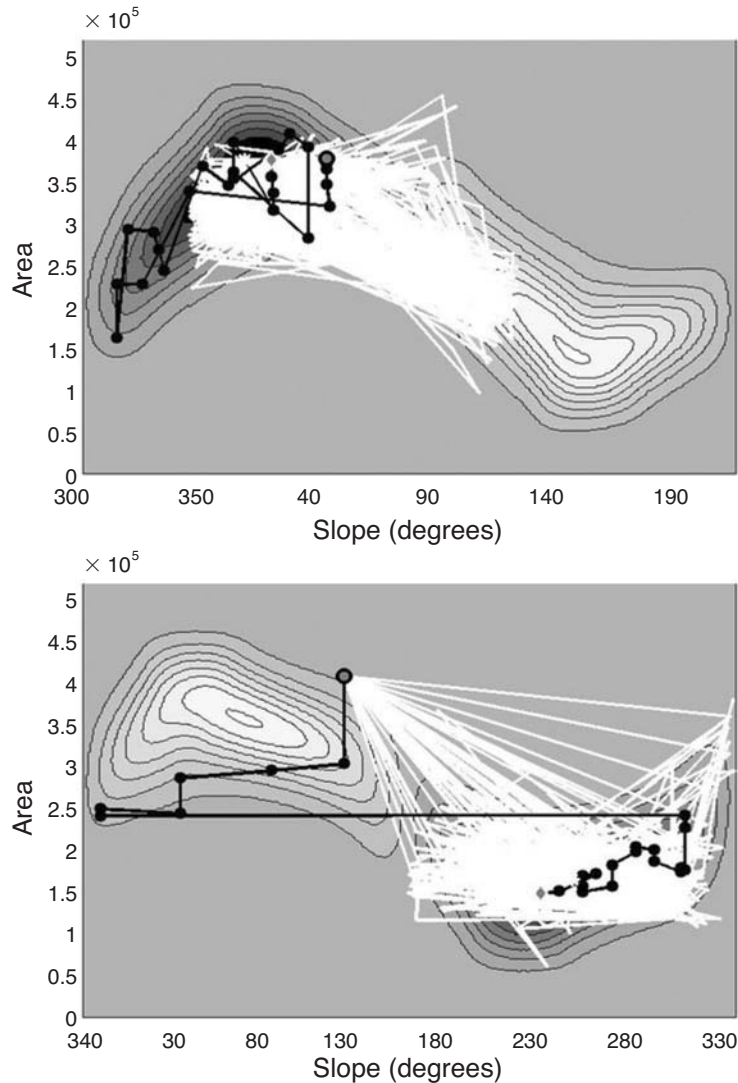


Figure 7. The 95% confidence regions (in white) simulated by the optimal learner for CS1-INI4 (top) and CS3-INI1 (bottom). Overlaid upon the confidence regions are examples of trajectories (in black) classified as violating gradient descent with respect to the optimal learner. The gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound. The optimal bounds are obscured by the confidence intervals and are located at the coordinates 30° , 3.75×10^5 pixels² (top), and 250° , 1.65×10^5 pixels² (bottom).

sion bound theory (Ashby, 1992a; Maddox & Ashby, 1993). The GLC assumes that observers always use a linear decision bound. The optimal-learning model that we added to the GLC assumed perfect memory for all of the stimuli that had been presented. In addition, after each stimulus presentation, the parameters of the decision bound were reestimated to those values that minimized the proportion of errors up to that point. Of course, since the optimal learner initially had no knowledge of the category parameters, its representation of the local gradient

information was imperfect, although this representation improved with experience.

The performance of the optimal learner was simulated for a total of 250 replications of 200 trials for all possible category structure-initial decision bound combinations (see the Appendix for details of all the simulations). A composite deviation index was computed for each of the 250 simulated trajectories, using the following procedure. First, an average simulated trajectory was obtained by computing the average values of the deci-

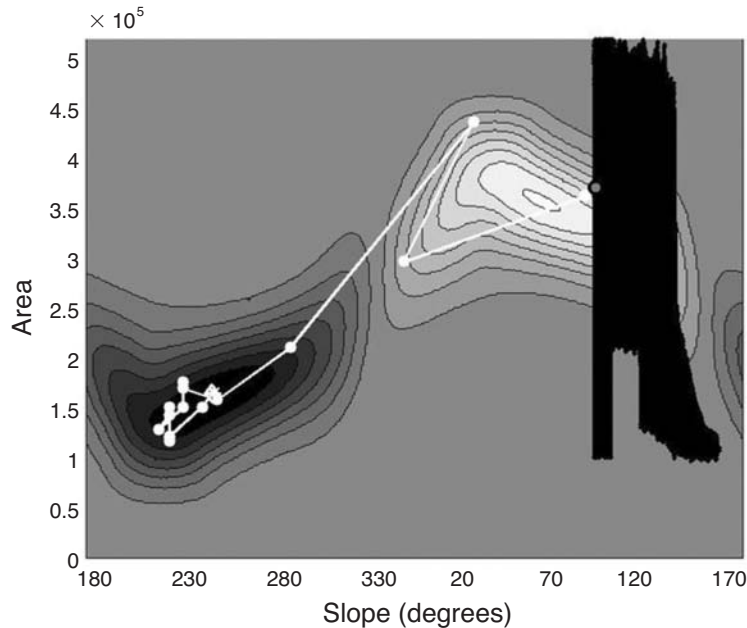


Figure 8. The 95% confidence region (in black) simulated by the adaptive learner for CS3–IN14. Overlaid upon the confidence region is a trajectory (in white) classified as violating gradient descent with respect to the adaptive learner. The gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound. Both the final and the optimal bounds are obscured by the trajectory and are located at the coordinates 250° , 1.65×10^5 pixels².

sion bound parameters on each trial across replications. Next, the deviation from the average trajectory (in Euclidean distance) was computed on a trial-by-trial basis for each replication.⁵ Finally, these deviations were standardized, and a composite index for each simulated trajectory was computed by summing the standardized scores for all 200 points. The 250 trajectories were then rank ordered according to their deviation index, and the most deviant 5% of the simulated trajectories were deleted. This procedure resulted in the construction of a 95% confidence region determined by finding the envelope of the remaining trajectories (i.e., the 95% that were most typical).

Representative confidence regions for two of the category structures are plotted in white in Figure 7. For exposition, the confidence regions in Figure 7 were generated by plotting the 95% most typical simulated trajectories. There are several points worth noting about the confidence regions. First, in both category structures, the confidence regions include portions of the parameter space that have a probability of error less than .50. Second, in the top panel, where the initial decision bound was near the optimal decision bound, the confidence region also includes points with a probability of error greater than .50. Third, there was clearly a considerable amount of variability in the simulated trajectories, as indicated by the spread of the confidence regions. These observations support the claim that even an optimal gradient descent

learning algorithm could produce trajectories that depart substantially from the trajectory predicted by gradient descent over the theoretical (i.e., population-based) probability-of-error surface.

To determine whether the empirical trajectories were consistent with sample-based gradient descent, we compared each of the empirical trajectories with the appropriate confidence region. Specifically, a trajectory was classified as violating sample-based gradient descent if more than 5% of the movements (i.e., changes in the state of the empirical trajectories) were to points outside of the confidence region. Analyzing the trajectories in this way revealed that 66% of the observers violated gradient descent. On average, trajectories that were classified as violating gradient descent produced 25 movements during the course of learning. It could be argued that given this small number, a criterion of 5% is too liberal. However, increasing this criterion to 20% still results in almost half (42%) of the trajectories being classified as violating gradient descent. Two examples of trajectories that violated gradient descent are overlaid in black on the confidence regions in Figure 7. Although not absolute, this analysis indicates that gradient descent, by itself, does not provide a complete description of these data.

In addition to investigating the confidence regions predicted by the optimal learner model, it is also worthwhile to examine the confidence regions predicted by

other popular gradient descent algorithms. As a result, we investigated the predictions of a version of the GLC that learned via a least mean square (LMS) algorithm, and we also investigated the predictions of Kruschke's (1992) ALCOVE (attention learning covering map) model. As with the optimal learner, confidence regions were computed for both models and compared with the empirical trajectories. The goal of these analyses was not to test these models rigorously but, rather, to explore the predictions of several different gradient-descent-based learning algorithms.

Adaptive learner. The adaptive learner is a simple connectionist network version of the GLC. It differs from the optimal learner only in its learning assumptions (and it had no internal noise). The adaptive learner used a popular connectionist gradient descent learning algorithm—namely, the delta rule (Widrow & Hoff, 1960). Details are given in the appendix.

The simulations are based on 1,000 replications using a single category structure (CS3) and initial bound (INI3). The confidence region was computed using the same procedure as that for the optimal learner and is plotted in Figure 8 (this time in black). Inspection of the confidence region illustrates that this simple model did not learn the categories within the allotted 200 trials—that is, the trajectories did not terminate at the global minimum of the error surface. However, simulation runs with an increased number of trials demonstrated that the adaptive learner does eventually reach the global minimum. Furthermore, increasing the learning rate (i.e., η in the Appendix) above the value used in these simulations resulted in convergence to the optimal decision bound in fewer trials, but at the expense of increased variability in the confidence regions. In addition, consistent with the optimal learner, the confidence regions were initially fairly wide and narrow with time—indicating that the adaptive learner builds a more detailed representation of the error surface with experience.

To test whether the empirical trajectories were consistent with the trajectories of the adaptive learner, we replicated the analysis performed with the optimal learner, but only for empirical trajectories with the CS3–INI3 combination ($n = 5$). Not surprisingly, all of the empirical trajectories were inconsistent with the adaptive learner. This result was expected given that the adaptive learner did not learn the categories. Interestingly, not all of the observed violations were driven primarily by decision bounds near optimal. Consider, for example, the white trajectory overlaid upon the adaptive learner confidence region in Figure 8. In direct contrast to the adaptive learner, this trajectory proceeds in the opposite direction along the error surface.

ALCOVE. ALCOVE is an exemplar-similarity-based connectionist model of category learning. Briefly, it assumes that category decisions are made by computing the similarity of the stimulus to memory representations of all the previously seen exemplars (see Kruschke, 1992, for a full description). ALCOVE has four param-

Table 3
Parameter Values Used in the ALCOVE Simulations

c	ϕ	λ_w	λ_a
.75	4.5	.01	.005
4.5	4.5	.01	.005
4.5	4.5	.1	.005
.75	4.5	.1	.005
.75	4.5	.01	.16
4.5	4.5	.01	.16
4.5	4.5	.1	.16
.75	4.5	.1	.16

eters: c , ϕ , λ_w , and λ_a . The c parameter is a measure of the overall discriminability of the stimuli, ϕ specifies the consistency of responding, and the two learning rates, λ_w and λ_a , determine how quickly the exemplar category associations and attention weights are learned, respectively. ALCOVE is a popular model of category learning that assumes that the weights between layers in the network are updated by a process of gradient descent on error.

For the present application, two numerical values of c , λ_w , and λ_a were used, and the value of ϕ was held constant—resulting in eight unique parameter combinations (see Table 3). These parameter values have all been used in past applications of ALCOVE (e.g., Nosofsky & Kruschke, 2001). Each of the eight parameter combinations was used to generate simulated trajectories for all four category structures (using a single initial bound, INI3). Each of the 32 possible parameter–category–structure combinations were replicated 200 times. Confidence regions, which were generated using a procedure identical to that for the optimal and the adaptive learners, are shown in white in Figure 9.

The ALCOVE trajectories were tested as before by comparing the empirical trajectories ($n = 20$) with the appropriate confidence regions. Each trajectory was individually compared with all eight confidence regions generated by varying the parameters of ALCOVE. Using the criterion that gradient descent was violated if more than 5% of the movements for a trajectory were to points outside of the confidence region produced results consistent with those from the analysis of the optimal learner—that is, 65% of the empirical trajectories were classified as violating ALCOVE. However, unlike the optimal learner analysis, increasing the criterion to 20% did not produce a substantial decrease in the percentage of trajectories that were classified as violating gradient descent (i.e., 62.5%). Examples of trajectories that were consistent and inconsistent with the ALCOVE predictions are shown in black in the top and bottom panels of Figure 9, respectively.

A related question is whether the magnitude of the movements made by the observers is consistent with the magnitude of the movements made by the models considered in the last section. In particular, an interesting question is whether the magnitude of the average movement following incorrect trials for the simulated models and the human observers are comparable. In fact, this was not the case. The magnitude of the change in the de-

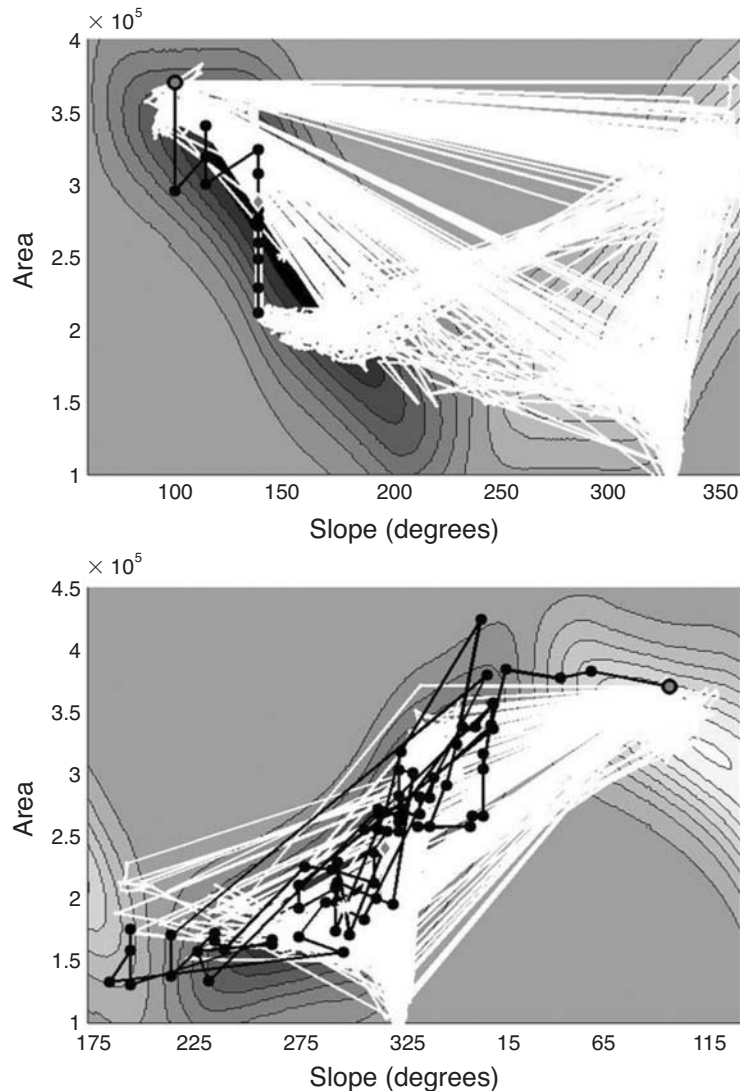


Figure 9. The 95% confidence region (in white) simulated by ALCOVE for CS2 (top) and CS4 (bottom). Overlaid upon the confidence regions are trajectories (in black) classified as not violating (top) and violating (bottom) gradient descent with respect to ALCOVE. The parameter values for both confidence intervals were $c = .75$, $\phi = 4.5$, $\lambda_w = .01$, and $\lambda_a = .005$. The gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound. The optimal bounds are obscured by the confidence intervals and are located at the coordinates $150^\circ, 2.51 \times 10^5$ pixels² (top), and $300^\circ, 1.93 \times 10^5$ pixels² (bottom).

cision bounds of the human observers was far greater ($M = 138.63$, $SE = 5.96$) than that for the optimal learner ($M = 52.65$, $SE = 14.58$), the adaptive learner ($M = 8.87$, $SE = 1.51$), and ALCOVE ($M = 13.55$, $SE = 3.49$).

Cooling

As another illustration of the type of question that can be investigated using dynamical trajectories, we will now focus on the degree to which these data are consistent with cooling—or a general decrease in step size with

experience. Although cooling is often incorporated in learning models, to our knowledge it has never been empirically tested. Because changes in the decision bound following correct trials were smaller in magnitude than those following incorrect trials, and since the number of correct trials increases with training, it follows that there was a general decrease in step size with experience. Thus, the question was not whether such a decrease occurred, but rather whether the decrease was greater than that predicted by gradient descent.




Figure 10. Average value of the cooling statistic, $C (\pm SEM)$, across trial blocks.

One way to answer this question is to compare the observed decreases in step size with the local gradient. Gradient descent algorithms predict that the step size on each trial will be proportional to the local gradient (see Equation 1). For example, let $S(n)$ denote the step size on trial n , and let $G(n)$ denote the gradient in the direction of steepest descent on trial n . Then gradient descent predicts that

$$S(n) = \gamma G(n),$$

for some constant γ or, in other words, that the ratio $S(n)/G(n)$ is a constant as n increases. For this reason, we defined the following cooling statistic:

$$C(n) = \gamma S(n) / \delta G(n),$$

where $G(n)$ was computed as the difference in the probability of error between the starting point on trial n and a step size of five units in slope–intercept space (see note 4) in the direction of the decision bound on trial $n + 1$.⁶ Note that $G(n)$ is negative any time an observer moves in a direction that increases the probability of error, positive if the probability of error was smaller on trial $n + 1$ than on trial n , and 0 if there was no change in accuracy. The constants γ and δ were set to the largest gradient possible (i.e., surface maximum – surface minimum = 0.6) and the largest step size possible (1,456 pixels; see note 4), respectively.

One prediction is that the mean of $C(n)$ across trials, \bar{C} will be positive if the participants are using gradient descent. Conversely, $\bar{C} < 0$ implies that the participants generally moved in a direction that increased the probability of error. The mean, \bar{C} , obtained by averaging across all the observers in all the category structures ($\bar{C} = 0.23$, $SD = 3.38$), was not significantly different from 0 [$t(624) = 1.73$, $p > .05$]. This implies that the partici-

pants, on average, did not generally move in directions that increased or decreased the probability of error. However, this analysis obscures differences throughout the course of learning.

To explore cooling in these data further, the change in $C(n)$ as a function of trial number (i.e., n) was analyzed. Recall that gradient descent predicts that $C(n)$ should not vary with n . If $C(n)$ decreases with n , then as observers approach the optimal bound, the step sizes are decreasing faster than the gradient, which indicates cooling. A finding that $C(n)$ increases with n indicates warming (i.e., the opposite of cooling).

The first step in this analysis was to obtain an aggregate cooling statistic by aligning each observer's data on the trial in which the stopping criterion was met and then averaging the absolute value of $C(n)$ across all the observers in all the category structures. This process initially resulted in 200 data points. Trials in which there were no movements were then omitted from the following analysis (35 trials total). The data were arbitrarily grouped into five blocks of 33 trials each (analyses in which different bin sizes were used produced similar results). The average value of C for each of the blocks is graphed in Figure 10. On average, C clearly decreased with experience. A one-way repeated measures analysis of variance was conducted on these data with the average value of C in each of the five blocks as the dependent variable. Consistent with a visual inspection of Figure 10, a significant decreasing linear trend was observed [$F(1,32) = 6.61$, $p < .05$]. Thus, these data appear to be inconsistent with simple gradient descent—that is, the participants decreased their movements at a faster rate than would be predicted by gradient descent. Note, however, that these violations of gradient descent

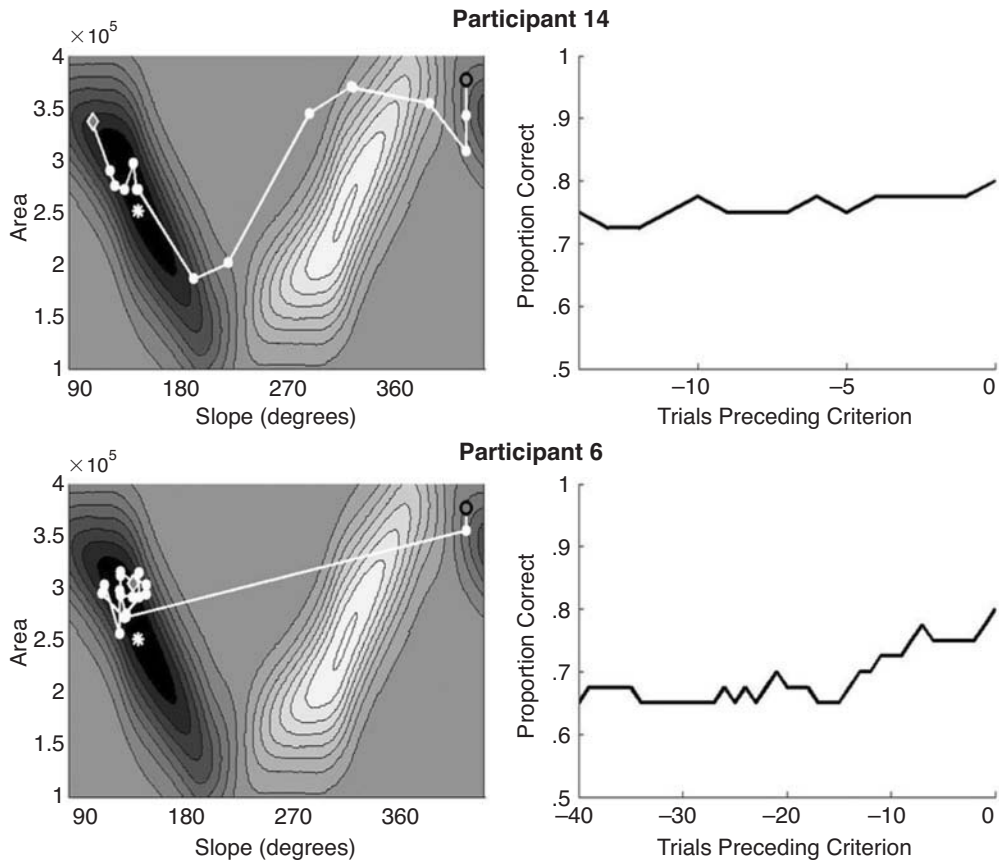


Figure 11. A comparison of empirical trajectories (left panels) and backward-learning curves (right panels) for 2 participants who learned CS2 (see the text for details). The 0 point on the learning curve graphs is the trial on which criterion was met. For the trajectories, the gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound.

affect only the speed with which the optimal bound is approached, not the path taken through the parameter space. As such, the existence of cooling could not be responsible for the frequent rejections of gradient descent described in the previous section.

In summary, the observers initially had changes in their decision bounds larger than those predicted by current models of category learning, and they decreased the magnitude of these changes at a faster than predicted rate.

Identifiability and Learning Curves

In addition to testing the assumptions of learning algorithms, the analysis of dynamical trajectories provides enough information to address the identifiability problems associated with learning curves. Recall the hypothetical category-learning experiment discussed earlier. It was shown that two very different models of the decision strategy were capable of producing nearly identical backward-learning curves. This argument can easily be extended to the present experiment if one considers individual differences in the decision strategy used by the observers. There are an infinite number of empirical trajectories that differ only in the path they follow to the optimal decision bound

(i.e., share only the initial and final decision bounds) but produce identical changes in accuracy across trials.

As an example, consider the learning curves and empirical trajectories from CS2 and CS3 in Figures 11 and 12, respectively. The left panels in both figures are the empirical trajectories of 2 observers, both with the same initial decision bound. The right panels in both figures are plots of the backward-learning curves for each of the illustrated trajectories (0 on the horizontal axis corresponds to the trial on which the stopping criterion was met, and negative trial numbers refer to preceding trials). It is clear from the abscissas of the backward-learning curves plotted in Figure 11 that Participant 14 reached criterion sooner than did Participant 6. In addition, changes in proportion correct were much more frequent for Participant 6 than for Participant 14; thus, one inference from these backward-learning curves might be that Participant 6 cycled through a series of extremely different decision bounds. However, an investigation of the empirical trajectories reveals just the opposite. Participant 6 used a series of similar decision bounds at or near the optimal bound, whereas Participant 14 tried many different decision bounds, few of which were optimal.

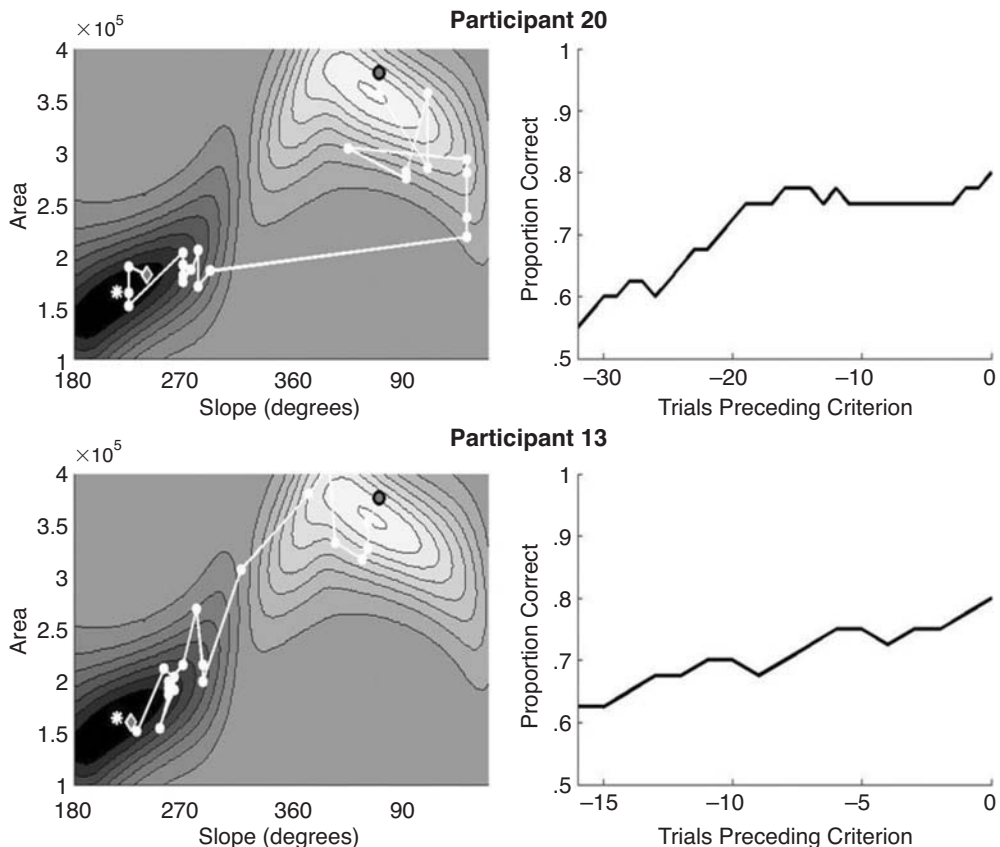


Figure 12. A comparison of empirical trajectories (left panels) and backward-learning curves (right panels) for 2 participants who learned CS3 (see the text for details). The 0 point on the learning curve graphs is the trial on which criterion was met. For the trajectories, the gray circle (outlined in black) marks the initial bound, the gray diamond marks the final bound, and the white asterisk marks the optimal bound.

Figure 12 demonstrates a different situation. Clearly, the two backward-learning curves are not identical, but they are similar in the magnitude of the increase in accuracy and the general shape of the curves. Interestingly, these two, qualitatively similar, backward-learning curves were estimated from two dissimilar learning trajectories. Both participants started at the same point on the error surface and eventually reached the optimal bound, but the majority of the decision bounds used by Participant 20 were almost 180° off from the optimal bound. In sum, the examples in Figures 11 and 12 illustrate two different cases in which it would be difficult to draw accurate inferences about the evolution of the observer's decision strategy if backward-learning curves were the sole analytic tool.

Conclusions

The preceding analyses illustrate the diversity of questions that can be addressed by investigating dynamical trajectories. Primary among these was empirically testing the assumption of gradient descent in category learning. A set of simulation analyses using three different gradient-descent-based models revealed that gradient

descent did not provide an adequate description of our data—that is, the majority of the observable trajectories violated gradient descent. Furthermore, we found strong evidence of cooling—that is, changes in decision strategy decreased at a faster rate than predicted by gradient descent. Finally, we provided an empirical demonstration of the identifiability problems inherent in learning curves. Specifically, it was shown that learning curves are not reliable as a basis for inferring an observer's decision strategy.

EXPERIMENT 2

A potential criticism of Experiment 1 is that forcing the observers to use linear decision bounds throughout the experiment placed unreasonable constraints on their decision strategies and, therefore, masked the underlying process of interest. Perhaps the observers would have adopted more complex decision bounds (e.g., quadratic) to learn the category structures in Experiment 1 if the linear constraint had not been imposed. Thus, the goal of Experiment 2 was to validate the methodology in Exper-

iment 1 by using a more traditional category-learning paradigm.

Method

Participants

An additional 20 undergraduates participated in Experiment 2.

Stimuli and Apparatus

All the stimuli were presented as light blue circles on a dark blue background.

Procedure

The procedure was identical to that in Experiment 1, with the following exceptions. On a trial, a single stimulus was presented, the observer was instructed to make a category assignment by depressing one of two response keys (labeled "A" or "B") with the index finger, and trial-by-trial feedback was provided. The trials were observer paced with an upper time limit of 5 sec. If a response was not given in that time period, the observer was prompted to speed up his or her response, and that trial was discarded. The response-stimulus interval was 500 msec. In an effort to equate the observed difficulty across the two experiments, the stopping criterion was decreased to a performance level of 75% correct (from 80% in Experiment 1) over the last 40 trials to move on to the next category structure or exit the experiment.

Results and Discussion

Trials to Criterion

Each of the 20 participants attempted to learn four category structures. Two participants met criterion in only one of the category structures and were not considered in any of the analyses. Seven category structures (out of 72 total) in which the participants did not meet the stopping criterion were not considered in the following analyses. The median number of trials to reach the stopping criterion for each category structure is given in Table 4. None of these differences were statistically significant [$\chi^2(3) = 3.13$].

Modeling Analyses

The goal of this analysis was to test the hypothesis that the more constrained Experiment 1 paradigm did not significantly alter the decision strategy that would have been used in a more typical category-learning task. Accordingly, we fit a number of different decision bound models (see Ashby, 1992a, and Maddox & Ashby, 1993, for detailed descriptions of the models) to the last 40 trials of each category structure for each participant in both experiments.⁷ The following three decision bound models were fit to the responses of each observer.

The General Linear Classifier. This model assumes that the decision bound between each pair of categories is linear. The GLC has three parameters: the slope and the intercept of the linear bound and the variance of internal (perceptual and criterial) noise (i.e., σ^2). A special case of the GLC assumes that observers use the linear bound that maximizes accuracy (i.e., the solid bounds shown in Figure 4). This model has only one free parameter (σ^2) and is referred to as the *optimal classifier* (OC).

The General Quadratic Classifier (GQC). A natural extension of the GLC is to assume that the observer uses a quadratic, rather than a linear, decision bound. The GQC has six free parameters (five describing the form of the decision bound and σ^2).

All of the models were fit using an iterative maximum likelihood search procedure (Ashby, 1992a; Wickens, 1982), and the goodness-of-fit statistic was determined by computing the AIC value ($AIC = 2r - 2\ln L$, where r is the number of free parameters and L is the likelihood of the model given the data; Akaike, 1974). The AIC statistic penalizes a model for poor fit and for the number of free parameters. To find the best model among a set of competitors, one simply computes an AIC value for each model and then selects the model with the smallest AIC.

The percentage of data sets that were best accounted for by each model across both experiments, along with the average percentage of responses accounted for (RA) by the best-fitting model, is listed in Table 5. As was expected, the majority of the data from both experiments were best fit by a linear decision bound (i.e., GLC and OC), and importantly, this percentage was almost identical across the two experiments.⁸ Furthermore, the RA suggests that the models were successfully accounting for the observed categorization data, thereby strengthening the results of this analysis.

Conclusions

Experiment 2 served as a control to Experiment 1 in that the goal was to determine whether forcing the observers in Experiment 1 to adopt linear decision bounds fundamentally altered category learning. The results of Experiment 2 suggest that performance was well described by a linear decision bound, even when no constraints were placed on the form of the decision bound. Therefore, the results obtained from the Experiment 1 paradigm should not be considered an artifact but, rather, a valid method for studying category learning.

GENERAL DISCUSSION

Learning curves are arguably the most widely employed technique in the analysis of category-learning data, but this level of data analysis can obscure meaningful differences between models, even at the individual observer level. We argued that a dynamical systems approach to category learning enriches the data that can

Table 4
Median Trials to Criterion in the Four Category Structures of Experiment 2

Category Structure	Median	Interquartile Range
1	63	43.00
2	69	101.25
3	52	27.00
4	49	35.00

Table 5
Percentage of Data Sets That Were Best Fit by the GLC, the GQC, or the OC According to AIC

	Model			RA	
	GLC	GQC	OC	<i>M</i>	<i>SD</i>
Experiment 1	57	12	31	90	0.08
Experiment 2	60	12	28	81	0.10

Note—RA, percentage of responses accounted for.

be obtained and provides for analyses at a much finer level of detail than is permitted by learning curves. Experiment 1 introduced a novel paradigm for making the dynamical learning trajectories directly observable. This experimental technique was used to test the assumption of gradient descent in category learning and to illustrate the identifiability problems associated with learning curves. Experiment 2 established that the dynamical trajectory paradigm is a valid method for investigating category learning.

Analysis of the dynamical learning trajectories provided evidence against the hypothesis that categories are learned by a simple gradient descent process. The results from the qualitative inspection of the empirical trajectories and the analysis of cooling were not consistent with the assumption of gradient descent, but both analyses assumed that the observers had access to complete information about the theoretical error surface. Simulations from three different models that assumed gradient descent produced a more rigorous test, which also disconfirmed simple gradient descent.

Although gradient descent is apparently incompatible with our data, our results do not rule out a more complex learning process in which gradient descent is augmented with a simulated annealing process (Darken & Moody, 1992; Geman & Geman, 1984). Simulated annealing, which has its roots in statistical thermodynamics, would have the effect of shaking the dynamical trajectories as gradient descent forces them downhill. The result would produce jumps in the trajectories that could account for the larger-than-predicted step sizes and the frequent excursions outside of the gradient descent confidence regions.

Another possibility that our results cannot rule out is that gradient descent is augmented by an explicit rule-based process of hypothesis testing (Ashby et al., 1998). According to this idea, observers occasionally generate specific verbalizable hypotheses about category structure. Testing these hypotheses introduces large discrete jumps in the trajectories. Ashby et al. (1998) hypothesized that such a process could mimic simulated annealing. The dynamical trajectories estimated by Alfonso-Reese (1996) were consistent with this rule-based hypothesis.

A related issue is whether or not gradient descent may be more likely depending on the type of category-learning task. A recent distinction has been made between tasks that can be learned via some explicit hypothesis-testing procedure (i.e., rule-based tasks) and those tasks that require the observer to combine information from multiple

stimulus dimensions at some predecisional stage (i.e., information integration tasks; Ashby & Ell, 2001; Maddox & Ashby, 2004). Observers are able to learn both types of tasks but are rarely able to provide an accurate description of their decision strategy in information integration tasks. Two features of the present experiments suggest that these tasks are more similar to rule-based than to information integration tasks. First, although it is possible to attend to both the horizontal and the vertical positions of the stimuli, it is perhaps even more likely that these two sources of information were encoded as a single spatial position dimension. Second, clearly the participants had trial-by-trial access to the decision criterion in Experiment 1. However, it could be argued that they also had access to a very accurate, albeit noisy, representation of the decision criterion in Experiment 2. Thus, it may be the case that information integration tasks, but not rule-based tasks, are learned by gradient descent.

A comparison of the dynamical trajectories to the backward-learning curves provided an empirical demonstration of the identifiability problems with learning curves. It was shown that observers can have similar learning curves despite clear differences in the decision strategies employed during the course of learning. Therefore, it is difficult to make strong inferences about the mechanisms underlying learning when the only data available are learning curves.

Experiment 2 served as a control experiment to validate the technique used to obtain dynamical learning trajectories in Experiment 1. Specifically, the goal was to test whether constraining observers to use only linear decision bounds (Experiment 1) in some way altered the approach they would have taken to learn the categories under conditions more representative of typical category-learning tasks. We found that the decision strategy of observers in a two-alternative, forced choice category-learning task was best described by linear decision bounds in 88% of the data sets, thereby supporting the Experiment 1 paradigm.

The most obvious difficulty in utilizing the dynamical trajectory paradigm is developing a method for estimating the parameters that define the state space and the trajectories. Our solution to this problem was to define the state space in terms of the parameters that specify the decision bounds and devise an experiment in which the decision bounds—and therefore, the parameters—were directly observable on a trial-by-trial basis. For simplicity, we chose to use category structures for which a linear decision strategy was optimal, but this paradigm is not limited to this choice. For example, one could devise an experiment that used a more complicated conjunctive or quadratic decision strategy.

In the present study, the initial decision bounds were arbitrary. The only minor constraint that we added was to begin the participants at a variety of locations in the parameter space. Clearly, the starting position in this paradigm is important. Consider 2 participants, 1 of whom

starts with a bound that has a probability of error less than .50 and is near the minimum on the error surface (i.e., the optimal decision bound). The other participant begins with a decision bound that is far from optimal and has a probability of error greater than .50. Intuitively, one would expect this difference to have an effect on the participants' trajectories and overall performance in the task. Therefore, a potential next step with this paradigm is to test rigorously the influence of the initial decision bound on category learning.

The category structures could also be constructed so that stronger tests of gradient descent are possible. In particular, learning algorithms assuming gradient descent are known to be susceptible to local minima (e.g., Busemeyer, Myung, & McDaniel, 1993; Rumelhart et al., 1986; Widrow & Hoff, 1960). In the context of the present experiments, a local minimum would be a decision bound on the error surface where the gradient is zero in every direction and for which any reasonably small change in the decision bound results in an increase in the probability of error. With the normally distributed categories used in this article, no such local minima exist. However, it is straightforward to construct categories for which the resulting error surface is characterized by one or more local minima. In such cases, gradient descent learning algorithms predict that any trajectory passing through a local minimum will become trapped there, and asymptotic accuracy will, therefore, be suboptimal (White, 1989). In resource allocation tasks, it has been shown that observers are susceptible to local minima, in that information is sometimes combined in a manner that produces suboptimal performance even if the optimal performance criterion is known (Busemeyer et al., 1986). The use of the dynamical trajectory paradigm presented here could provide an extension of this research to category learning.

Experiments in which the decision criterion is observable on a trial-by-trial basis (Busemeyer & Myung, 1992; Kubovy & Healy, 1977) and in which dynamical learning trajectories are estimated (Alfonso-Reese, 1996; Minda & Smith, 2001; Smith & Minda, 1998, 2000) have been used to distinguish between competing models of the same psychological process. The estimation of dynamical trajectories could also serve to help distinguish between component systems within a multiple systems framework. Arguments for multiple processing systems that mediate learning have been made in such diverse fields as reasoning (Sloman, 1996), motor learning (Keele, Ivry, Mayr, Hazeltine, & Heuer, 2003; Willingham, Nissen, & Bullemer, 1989), discrimination learning (Kendler & Kendler, 1962), function learning (Hayes & Broadbent, 1988), and category learning (Ashby et al., 1998; Ashby & Ell, 2001; Brooks, 1978; Erickson & Kruschke, 1998; Minda & Smith, 2001; Pickering, 1997). Consequently, there is a growing need for analytic techniques that are capable of assessing the relative contribution of these component systems to learning.

Recall the hypothetical category-learning experiment discussed in the introduction, in which it was shown that backward-learning curves were not a powerful method for identifying the use of different learning algorithms. When the stimuli were constructed from continuous-valued dimensions, there was an observable signature in the backward-learning curves indicating a difference between the rule-based and the procedural-learning algorithms. However, when the stimuli were constructed from binary-valued dimensions, the signature was no longer present. This hypothetical example maps quite well onto the issue of multiple versus single systems of learning and memory. For example, it might be possible to use dynamical trajectories to analyze the differences between different systems that each contribute to learning, which would otherwise be undetectable if only learning curves were available.

SUMMARY AND CONCLUSIONS

A great deal of information is lost by examining only learning curves. A more powerful alternative is to examine trial-by-trial learning in the form of dynamical learning trajectories. The primary challenge in utilizing this approach is estimating the parameters that specify the decision bounds. The paradigm presented in this article solved this problem, thereby allowing for a much more detailed account of category learning. The use of dynamical learning trajectories allows not only asymptotic performance to be examined, but also the analysis of trial-by-trial information. We then used the dynamical trajectory paradigm to investigate the often assumed, but rarely tested, assumption of gradient descent. Our data indicate that gradient descent, at least in the strictest sense, does not accurately describe the dynamics of category learning.

REFERENCES

- AKAIKE, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**, 716-723.
- ALFONSO-REESE, L. (1996). *Dynamics of category learning*. Unpublished doctoral dissertation, University of California, Santa Barbara.
- ASHBY, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-483). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G. (1992b). Multivariate probability distributions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1-34). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN, A. U., & WALDRON, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, **105**, 442-481.
- ASHBY, F. G., & ELL, S. W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences*, **5**, 204-210.
- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- ASHBY, F. G., & MADDOX, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, **37**, 372-400.
- BRAINARD, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, **10**, 433-436.
- BROOKS, L. (1978). Nonanalytic concept formation and memory for in-

- stances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.
- BUSEMEYER, J. R., & MYUNG, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory and human performance. *Journal of Experimental Psychology: General*, **121**, 177-194.
- BUSEMEYER, J. R., MYUNG, I. J., & MCDANIEL, M. A. (1993). Cue competition effects: Theoretical implications for adaptive network learning models. *Psychological Science*, **4**, 196-202.
- BUSEMEYER, J. R., SWENSON, K. N., & LAZARTE, A. (1986). An adaptive approach to resource allocation. *Organizational Behavior & Human Decision Processes*, **38**, 318-341.
- DARKEN, C., & MOODY, J. E. (1992). Towards faster stochastic gradient search. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *Advances in neural information processing systems* (pp. 1009-1016). San Mateo, CA: Morgan Kaufman.
- ERICKSON, M. A., & KRUSCHKE, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, **127**, 107-140.
- ESTES, W. K. (1993). Models of categorization and category learning. In G. V. Nakamura & D. L. Medin (Eds.), *Categorization by humans and machines* (pp. 15-56). San Diego: Academic Press.
- ESTES, W. K. (1994). *Classification and cognition*. Oxford: Oxford University Press.
- ESTES, W. K., CAMPBELL, J. A., HATSPOULOS, N., & HURWITZ, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 556-571.
- GEMAN, S., & GEMAN, D. (1984). Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **6**, 721-741.
- GLUCK, M. A., & BOWER, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227-247.
- HAYES, N. A., & BROADBENT, D. (1988). Two modes of learning for interactive tasks. *Cognition*, **28**, 249-276.
- HAYKIN, S. (1994). *Neural networks*. Englewood Cliffs, NJ: Macmillan.
- HURWITZ, J. B. (1990). *A hidden-pattern unit model of category learning*. Unpublished doctoral dissertation, Harvard University.
- KEELE, S. W., IVRY, R. B., MAYR, U., HAZELTINE, E., & HEUER, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological Review*, **110**, 316-339.
- KENDLER, T. S., & KENDLER, H. H. (1962). Inferential behavior in children as a function of age and subgoal constancy. *Journal of Experimental Psychology*, **64**, 460-466.
- KRUSCHKE, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, **99**, 22-44.
- KUBOVY, M., & HEALY, A. F. (1977). The decision rule in probabilistic categorization: What it is and how it is learned. *Journal of Experimental Psychology: General*, **106**, 427-446.
- LEE, W., & JANKE, M. (1964). Categorizing externally distributed stimulus samples for three continua. *Journal of Experimental Psychology*, **68**, 376-382.
- LEE, W., & JANKE, M. (1965). Categorizing externally distributed stimulus samples for unequal molar probabilities. *Psychological Reports*, **17**, 79-90.
- MADDOX, W. T., & ASHBY, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, **53**, 49-70.
- MADDOX, W. T., & ASHBY, F. G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behavioral Processes*, **66**, 309-332.
- MINDA, J. P., & SMITH, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 775-799.
- MINSKY, M. L., & PAPERT, S. A. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- NOSOFSKY, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, **115**, 39-57.
- NOSOFSKY, R. M., & KRUSCHKE, J. K. (2001). Single-system models and interference in category learning: Commentary on Waldron and Ashby (2001). *Psychonomic Bulletin & Review*, **9**, 169-174.
- NOSOFSKY, R. M., KRUSCHKE, J. K., & MCKINLEY, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 211-233.
- PELLI, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, **10**, 437-442.
- PICKERING, A. D. (1997). New approaches to the study of amnesic patients: What can a neurofunctional philosophy and neural network methods offer? In A. R. Mayes & J. J. Downes (Eds.), *Theories of organic amnesia* (pp. 255-300). Hove, U.K.: Psychology Press.
- ROSENBLATT, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, **65**, 386-408.
- RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning representations by back-propagating errors. *Nature*, **323**, 533-535.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461-464.
- SHEPARD, R. N., HOVLAND, C. L., & JENKINS, H. M. (1961). Learning and memorization of classification. *Psychological Monographs*, **75**(13, Whole No. 517).
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- SMITH, J. D., & MINDA, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 1411-1436.
- SMITH, J. D., & MINDA, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 3-27.
- WHITE, H. (1989). Some asymptotic results from learning in single hidden-unit feedforward network models. *Journal of the American Statistical Association*, **84**, 1003-1013.
- WICKENS, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco: Freeman.
- WIDROW, G., & HOFF, M. E. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electric Show & Convention, Convention Record*, **4**, 96-104.
- WILLINGHAM, D. B., NISSEN, M. J., & BULLEMER, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1047-1060.

NOTES

1. The rule-based model was the verbal component of the multiple system category-learning model, COVIS, developed by Ashby and colleagues (Ashby et al., 1998). The following are the parameters (and their numerical values) that were used in the simulations: increment/decrement in the salience of a dimension following correct/incorrect responses (.04/.01), $\lambda = 12$, $\gamma = 10$, initial learning rate = 0.30, momentum term = 0.93, and decay rate = 42. The initial salience and adjusted salience on each dimension were set to 0.33. A perceptron was used for the procedural-learning model with a learning rate of 0.05 (Minsky & Papert, 1969; Rosenblatt, 1958). All the parameters were crudely adjusted to generate the backward-learning curves for the stimuli defined on continuous-valued dimensions. The same parameter values were used to estimate the backward-learning curves for the binary-valued dimensions. For each model, a simulation consisted of 250 replications of 400 trials (eight 50-trial simulated blocks).

2. A backward-learning curve was generated separately for each replication, using the following procedure. The 1st simulated block (of 8 total) in which the criterion of 10 consecutive correct responses was met was renumbered as block 0. The simulated block immediately preceding the criterion block was renumbered as block -1 (the next block as -2, etc.), and the simulated block immediately following the criterion block was renumbered as block 1 (the next block as 2, etc.). This

renumbering procedure resulted in a total of 15 possible blocks, because criterion could have been met in any of the 8 simulated blocks (e.g., criterion could have been met in block 1 in one replication and block 8 in another replication). The next step was to align all 250 replications on the criterion block. In practice, several of the extreme blocks contained few data points; thus, any block that contained data from fewer than 5% of the replications was excluded (e.g., blocks -6 and -7 in Figure 1A). Finally, the average backward-learning curves in Figure 1 were generated by averaging across replications.

3. The intercept was defined as the point at which the decision bound intersected the perimeter of the monitor. The monitor intercept, rather than the intercept on the ordinate, was used because the latter is undefined for some values of the slope (e.g., 90°).

4. Although the slope and the area of the decision bound could be adjusted independently by the observer, there is a correlation between these two parameters with respect to the error surface. However, this correlation is inconsequential, since the dynamical trajectories were not estimated but, rather, directly observable.

5. All changes in the decision bounds were computed in Euclidean distance, using the original coordinates of the decision bounds: slope (in degrees) and monitor intercept (the point at which the bound intersected the perimeter of the monitor) of the decision bounds. The monitor intercept ranged from 0 to 2,912 pixels, where 0 defined the lower left corner of the monitor, and increased in the counterclockwise direction (e.g., monitor intercept = 1,456 defined the upper right corner). For the purpose of the distance computations, it was assumed that the maximum change in the monitor intercept was 1,456 pixels. Although the monitor intercept ranges from 0 to 2,912 pixels, 0 and 2,912 are the

same point (the lower left corner of the monitor). For example, a movement from 20 on trial n to 2,900 on trial $n + 1$ probably did not result from a change of 2,880 pixels, but rather 32 pixels. This assumption was also applied to changes in the slope (0° and 360° are the same point). All distances were computed using these corrected values. One way to address this difficulty in the future would be to constrain the magnitude of change on the slope and the monitor intercept dimensions to be 180° and 1,456, respectively.

6. The choice of step size to compute the local gradient information was arbitrary. The only constraint was that this step size should not be so large as to obscure significant changes in the local gradient information. Several step sizes were used in this computation (0.1, 1, 3, 5, and 10), and no qualitative differences were found.

7. Although the observers in Experiment 1 were constrained to use a linear decision bound on a trial-by-trial basis, it is possible that a linear bound would not provide a good account of performance across the last 40 trials. For example, consider an observer who applies a series of decision bounds in the Experiment 1 paradigm such that each new bound is merely an arbitrarily small rotation of the previous one. In this case, the decision bound estimated across 40 trials would more closely resemble a quadratic, rather than a linear, function. For this reason, the decision bound models were also fit to the last 40 trials of each category structure for each participant in Experiment 1.

8. AIC tends to favor more complex models (e.g., the GQC). Repeating the analysis with a statistic that increased the penalty for complex models ($BIC = r \ln N - 2 \ln L$, where N is the sample size; Schwarz, 1978) indicated that 97% of the data sets in both experiments were best accounted for by a model assuming a linear decision bound.

APPENDIX
Details of the Simulations

The performance of all the models was simulated over a number of replications (optimal learner, 250; adaptive learner, 1,000; and ALCOVE, 200) of 200 trials, using either all possible CS \times INI combinations (optimal learner), the CS3–INI3 combination (adaptive learner), or all CS \times INI3 combinations (ALCOVE). The stimuli used for the simulations were identical to those presented to the participants, and a different random presentation order was used in each replication. For all simulations, it was assumed that the decision bound changes only after incorrect trials.

Optimal Learner

The optimal learner was based on the GLC (Ashby, 1992a; Maddox & Ashby, 1993). The GLC has three free parameters (the slope and the intercept of the linear bound and the variance of internal noise, σ^2). The slope and the intercept of the linear decision bound were recomputed on a trial-by-trial basis so that, on any trial, the optimal learner used the bound that minimized the proportion of errors on all preceding trials. The optimal learner was an ideal observer, in the sense that it had perfect memory for all the stimuli it encountered. On the basis of previous research, σ^2 was fixed at a low value (i.e., 10) to reflect the low amount of internal noise that would be expected given the features of Experiment 1 (e.g., the decision criterion was directly observable, the display was high contrast, and stimulus duration was observer paced).

Adaptive Learner

The adaptive learner model was a three-layer connectionist network that instantiated the GLC. The three nodes in the input layer encoded the horizontal and the vertical positions of the stimulus and a constant bias term (i.e., the intercept of the linear decision bound), respectively. The single hidden unit computed the discriminant value, and the single output unit converted this to a response (i.e., to $a + 1$ or -1). The weight on the path between the hidden unit and the output unit was always fixed at 1. The three weights on the paths between the input units and the hidden unit were updated after every trial according to the LMS rule (Widrow & Hoff, 1960). According to this algorithm, following trial n , the weight between input unit i and the hidden unit is changed by

$$\Delta w_i = \eta[h^*(x_1, x_2) - h(x_1, x_2)]x_i,$$

where η is a learning rate parameter (set to 0.7) and $h(x_1, x_2)$ and $h^*(x_1, x_2)$ are the values of the current and optimal linear discriminant functions, respectively, for the stimulus pair (x_1, x_2) . When the weight on the bias term was updated, x_i was set equal to 1.

It was necessary to impose an additional constraint on the response-mapping procedure, because decision bounds that are a $\theta = 180^\circ$ rotation of each other will result in discriminant values of identical sign. Specifically, negative values of $h(x_1, x_2)$ were mapped to Category A responses for $0 < \theta \leq 180^\circ$ and to Category B responses for $180^\circ < \theta \leq 360^\circ$. Similarly, positive values of $h(x_1, x_2)$ were mapped to Category B responses for $0 < \theta \leq 180^\circ$ and Category A responses for $180^\circ < \theta \leq 360^\circ$. Furthermore, it was assumed that on a single trial, the change in $\theta \leq 180^\circ$.

ALCOVE

For the ALCOVE simulations, the number and position of the exemplar nodes in the hidden layer were updated on a trial-by-trial basis. ALCOVE can be viewed as a learning version of Nosofsky's generalized context model (GCM; Nosofsky, 1986). The GCM makes no mention of decision bounds, although in the present experiments, its contour of equivocality (the set of all stimuli for which the probability of responding A equals the probability of responding B) agrees asymptotically with the optimal linear bound (Ashby & Maddox, 1993). Therefore, we derived trajectories from ALCOVE by using the parameter mappings under which this equivalence holds (i.e., from Ashby & Maddox, 1993).
