


Article

The pigeon as a machine: Complex category structures can be acquired by a simple associative model



Brandon M. Turner, Edward A. Wasserman

turner.826@gmail.com

Highlights

Pigeons can solve an exceptionally broad range of visual categorization tasks

Some of these tasks seem to require advanced cognitive and attentional processes

Yet, computational modeling indicates that pigeons don't deploy such complex processes

A simple associative mechanism may be sufficient to account for the pigeon's success

Article

The pigeon as a machine: Complex category structures can be acquired by a simple associative model

Brandon M. Turner^{1,3,*} and Edward A. Wasserman²

SUMMARY

Never known for its smarts, the pigeon has proven to be a prodigious classifier of complex visual stimuli. What explains its surprising success? Does it possess elaborate executive functions akin to those deployed by humans? Or does it effectively deploy an unheralded, but powerful associative learning mechanism? In a series of experiments, we first confirm that pigeons can learn a variety of category structures – some devised to foil the use of advanced cognitive processes. We then contrive a simple associative learning model to see how effectively the model learns the same tasks given to pigeons. The close fit of the associative model to pigeons' categorization behavior provides unprecedented support for associative learning as a viable mechanism for mastering complex category structures and for the pigeon's using this mechanism to adapt to a rich visual world. This model will help guide future neuroscientific research into the biological substrates of visual cognition.

INTRODUCTION

The pigeon has never been celebrated for its smarts. Its intelligence has long been belittled, even by one of its most avid aficionados, the ethologist Charles Otis Whitman, who wrote that:

Whatever the pigeon instinct-mind contains, it is safe to say that intelligence is hardly more than a grain hidden in bushels of instinct, and one may search more than a day and not find it¹ (p.158).

Subsequently, the behaviorist B. F. Skinner joined the disparaging chorus in justifying his deployment of pigeons as the “brains” of an experimental guidance system to direct ballistic missiles to potential WWII military targets.² Skinner claimed to have chosen pigeons for his highly unorthodox project:

... not because the pigeon is an intelligent bird, but because it is a practical one and can be made into a *machine*, from all practical points of view³ (p. 851, italics added).

These negative appraisals notwithstanding, the pigeon has proven to be a particularly adept classifier of highly complex and varied visual stimuli. These stimuli range from photographs of everyday objects such as flowers, trees, dogs, ducks, keys, pens, phones, and shoes⁴ to medical images of human breast tissue⁵ and human heart muscle⁶ to printed numerals and letters of the alphabet⁷ and to artworks painted by the likes of Monet, Cezanne, and Renoir.⁸ The pigeon's excellent category learning and transfer performance clearly merit the label “conceptual” according to most theorists^{9–14} and suggest that we should accord pigeons and other birds far greater respect than they customarily receive.^{15,16}

Although it is altogether appropriate for us to acknowledge the pigeon's impressive categorization prowess, a far more profound and significant scientific question remains unanswered: Just how is the pigeon able to exhibit such noteworthy conceptual behavior?

Here, two viable interpretive paths present themselves. First, the pigeon may deploy some or all of the various executive functions commonly credited to humans and nonhuman primates, particularly selective attention and explicit rule use. Second, the pigeon may rely solely on associative learning, as it may possess none of those advanced executive functions.

How might we go about deciding between these interpretive possibilities? Because executive functions are generally deemed to be mediated by the prefrontal cortex (PFC), we might ask whether pigeons possess this or some other neuroanatomical structure – possibly the nidopallium collaterale (NCL) – that can perform the same cognitive functions. In this connection, our understanding of the avian brain and its relation to the mammalian neocortex is currently under intense investigation and substantial reanalysis.¹⁷ So, too, are the specific roles played

¹The Ohio State University, Columbus, OH, USA

²The University of Iowa, Iowa City, IA, USA

³Lead contact

*Correspondence: turner.826@gmail.com
<https://doi.org/10.1016/j.isci.2023.107998>



by numerous avian brain structures in perceiving and categorizing visual stimuli.¹⁸ Given the present state of flux in avian neurobiology, perhaps the second route should be given provisional priority.

Of course, associative learning is frequently presumed to be far too primitive and rigid to provide an adequate account of complex visual categorization. This pessimistic perspective is likely to be based on the sophistication and flexibility usually ascribed to human behavior and cognition. Nevertheless, the pigeon might represent a striking and enlightening counterexample.

The point of departure for our article concerns the unique relevance of the pigeon to a prominent neurocomputational theory of categorization: COVIS¹⁹ (COmpetition between Verbal and Implicit Systems). This theory hypothesizes that two distinctly different categorization mechanisms have evolved: (1) an implicit, procedural, or associative learning process that is slow and unanalytical and dependent on the basal ganglia, plus (2) an explicit, declarative, or verbal learning process that is fast and analytical and dependent on the PFC.

The cognitive accomplishments of the explicit process have largely taken center stage in the voluminous categorization literature, with far less attention being paid to the implicit process. This relative neglect may have contributed to a puzzling paradox²⁰: namely, that the power of associative processes in humans and nonhuman animals has been underestimated, whereas the power of associative mechanisms in the realm of artificial intelligence has been widely celebrated.^{21,22}

Nonetheless, the neurobiological foundation of the pigeon's categorization prowess may be based on an evolutionarily ancient and ostensibly elementary associative learning mechanism that progressively links behavioral responses to limited portions of perceptual space.^{11,23} Humans, while also possessing and ably deploying the same associative learning mechanism, may be able to exploit a far more elaborated declarative system which permits our more quickly and proficiently mastering an even wider range of categorization tasks by engaging stimulus analysis, selective attention, and rule formation.²³

The basic associative learning mechanism that is postulated for the pigeon has been lauded for its "simplicity, parsimony, breadth, and power"²³ (p. 2365). More intriguingly, it has been further speculated that: "A unitary, nonanalytic system might also be especially adept at learning complex or non-linear category-decision boundaries that would defeat a dimensionally aligned, rule-based system"²³ (p. 2365)."

A key piece of evidence on this matter concerns the dramatic disparity between humans' and pigeons' learning of Rule-Based (RB) and Information Integration (II) tasks. Unlike most categorization tasks, RB and II tasks are carefully crafted from the same bidimensional stimulus space – commonly entailing black-and-white sinusoidal gratings with bars that vary in spatial frequency and line orientation. RB tasks cut the stimulus space either vertically or horizontally, yielding a pair of subtasks in which different distributions of either spatial frequency (but not line orientation) or line orientation (but not spatial frequency) are associated with two different categorization report responses. By contrast, II tasks cut the same stimulus space along either the positive or the negative diagonal, yielding a pair of subtasks in which different distributions of spatial frequency and line orientation are associated with two different categorization responses. Thus, RB tasks can be solved using simple unidimensional rules and should encourage analytic declarative processing, whereas II tasks require featural or dimensional integration and should generally encourage nonanalytic associative processing.^{19,24}

Past research has reported that, unlike humans – who learn RB tasks much more quickly than II tasks – pigeons learn RB and II tasks at equal speeds.^{11,23,25} These findings accord with the view that pigeons lack the declarative categorization system and can only deploy the associative categorization system.

Because of the theoretical importance of this profound disparity in interspecies category learning, we chose to further assess the reliability of pigeons' indifference to learning RB and II tasks. Deploying both within- and between-subject designs, we too found no difference in the speed of pigeons' learning.²⁶ We later posed still greater challenges to pigeons' associative learning system by requiring them to learn newly devised tasks that entailed complex non-linear category decision boundaries that had been contrived to thwart a dimensionally aligned, rule-based system.^{26,27} The pigeons' stunning success in learning these new tasks pose the daunting challenge of explaining just how an elementary associative learning system might acquire such demanding tasks. That is just what we hope to accomplish in the following sections.

There, we will show through extensive computational modeling that pigeons can learn a variety of novel and difficult categorization tasks by utilizing an exceptionally simple associative mechanism – one that connects regions of multidimensional space to distinctive behavioral responses by means of an error correction process. The four tasks include: RB and II tasks involving spatially separated elliptical stimulus distributions; RB and II tasks involving abutting hemispheric stimulus distributions; a novel concentric-rings task; and a novel sectioned-rings task. In each case, the successful modeling of pigeons' task learning was complemented by accurately modeling characteristic signatures of associative learning, further bolstering our proposal that the pigeon singularly proves that complex cognitive behavior can nonetheless be accomplished via simple associative learning. Finally, the nature of our associative model effectively took Skinner at his word – by treating the pigeon as if it were a *machine*.

RESULTS

Figure 1A illustrates our general approach to understanding the learning mechanism used by the pigeon. On each trial, we assume that the pigeon encodes the displayed stimulus and computes the association between the stimulus percept and a categorization response based on an associative network. Figure 1B illustrates that the encoding of the stimulus activates a limited region of the pigeon's representation of the entire stimulus space (i.e., the hidden layer), and Figure 1C illustrates how the pattern of activation creates evidence for a categorization response based on learned associations between previous stimuli and feedback from previous responses. Based on the relationship between the learned associations and the current stimulus, a response is made from the output layer (Figure 1D).

To assess the power of associative learning, we developed a computational model that uses only the principles of stimulus-response association through reinforcement learning and stimulus generalization.²⁸ Because the model involves just two mechanisms, if the model can




Figure 1. The pigeon as a machine

(A) Illustration of a pigeon performing a category learning task by assigning a Gabor patch to one of the two categories. We assume that the pigeon applies a purely associative learning mechanism to map the contents of the stimulus into a response through a computational model very similar to those found in machine learning.

(B) The mapping from the stimulus space to the hidden layer is shown, where yellow areas represent regions of the perceptual space that are more strongly activated by a given Gabor patch.

(C) In the hidden layer, the model assumes each point in the perceptual space is mapped to evidence for each of the category responses. The mapping to Category A is shown on the left (red) and the mapping to Category B is shown on the right (blue), where brightly colored areas designate stronger associations.

(D) The pigeon produces a response based on the strength of evidence for each of the two categories.

successfully capture the pigeon's behavior in a wide variety of categorization tasks, then we can interpret the success of the model as evidence that the pigeon deploys only these two mechanisms. In other words, if the machine-like model behaves as does a pigeon, then it serves as an appropriate analogue of the principles underlying the pigeon's category learning, which will help guide future neuroscientific research into the biological substrates of visual cognition.

We used four different categorization tasks involving unique structures that afforded a stringent test of the pigeon-as-machine metaphor: (1) a task involving rule-based or non-rule-based conditions with an abrupt change in the category structure once pigeons attained the initial learning criterion, (2) a task involving concurrently trained rule-based and non-rule-based conditions involving abutting, but nonoverlapping stimulus distributions, (3) a task consisting of stimulus distributions that entailed two concentric rings, and (4) a task consisting of stimulus distributions that involved the same two rings, but now cut into eight separate sections. Although the four tasks differed considerably from one another, they all tested the pigeon-as-machine metaphor either by placing pigeons into learning environments where rule-based responding could not be deployed or, if rule-based responding could be deployed, where strategic comparisons between categorization tasks were able to detect rule use.

RB/II transfer task

Data from the first task were reported in O'Donoghue et al.,²⁶ where they trained pigeons on either an RB task in which a *single* dimension was sufficient for category learning or an II task in which *both* stimulus dimensions were needed for category learning. Pigeons were trained on these tasks until they achieved 85% accuracy to both Categories A and B for two consecutive sessions. At this time, the categorization boundaries changed such that the importance of the dimensions in their contributions to the tasks switched. For example, if the categorization boundary in the RB task was based on frequency information in the first phase, then the categorization boundary was based on orientation information in the second phase. Figures 2A, 2E, and 2I illustrate the distribution of stimuli in the RB task for Pigeon 1 on Trials 2 (at the beginning of Phase 1), 255 (early in Phase 2), and 640 (at the end of Phase 2), respectively; Figures 2C, 2G, and 2K illustrate the distribution of stimuli in the II task for Pigeon 5 on Trials 2 (at the beginning of Phase 1), 335 (early in Phase 2), and 560 (at the end of Phase 2), respectively.

After the model was fit to each pigeon's data, we examined the learned associations predicted by the model based on the estimated parameters extracted from the data-fitting process. Figures 2B, 2F, and 2J illustrate the associations learned between the stimulus space and each category over time (i.e., Category A on the left, Category B on the right), where each row corresponds to the same moment in time shown in Figures 2A, 2E, and 2I. In each panel, brightly colored areas correspond to strong associations, whereas white areas correspond to weak associations, and the black ellipsoid designates the region in which stimuli from the relevant category occur (i.e., 95% of the stimuli from that category). In Figure 2B, the model has no previous associations to either category, so the representations illustrate the spread of association




Figure 2. Representations learned in the RB/II transfer task

(A, E, I) Example distributions of stimuli used in the RB task over time. In panels E, G, I, and K, transparency is used to illustrate the order of the training stimuli over time.

(B, F, J) Learned associations between the stimulus space and Category A (left) and Category B (right) by the model at three points in time corresponding to the trials shown in A, E, and I.

(C, G, K) Example distributions of stimuli used in the II task over time.

(D, H, L) Learned associations between the stimulus space and Category A (left) and Category B (right) by the model at three points in time corresponding to the trials shown in C, G, and K. In all panels, brightly colored areas indicate strong associations, whereas white areas indicate weak associations, and each black ellipsoid designates the region in which stimuli from the relevant category occur. RB: rule-based task, II: information integration task.

following a single example from each category. At the beginning of Phase 2, shown in [Figures 2E](#) and [2F](#), the model has correctly mapped the stimulus space to the corresponding category response in Phase 1, and it begins to learn the new mapping in Phase 2. By the end of Phase 2, [Figures 2I](#) and [2J](#) show that the model has remapped the stimulus space in accord with the new categorization task. [Figures 2D](#), [2H](#), and [2L](#) illustrate that a very similar associative learning process occurs for the model in the II task. Specifically, the model learns the correct category mapping by the end of Phase 1; by the end of training, the model forms the correct association in accord with the new categorization task arranged in Phase 2.

[Figure 3](#) shows pigeons' accuracy data from the transfer task along with the model fits. To reduce visual overlap, each panel corresponds to a different comparison: either across the RB and II tasks (i.e., [Figures 3A](#) and [3C](#)) or across the two phases within a task (i.e., [Figures 3B](#) and [3D](#)). In each panel, the data are shown as the solid-colored line with the correspondingly colored shaded regions reflecting the standard error of the mean; the model fits are shown as the black lines with the colored dots corresponding to each condition. There are several important results, each of which provides strong evidence that the pigeons learned the categories purely by association.

First, [Figure 3A](#) reveals that there is no appreciable difference between the learning curves for the RB and II tasks in Phase 1. Second, [Figures 3B](#) and [3D](#) illustrate that neither the RB task nor the II task shows any appreciable difference in the speed of learning between Phases 1 and 2. Third, [Figure 3C](#) reveals that in Phase 2, there is no appreciable difference in learning the two tasks. Fourth, a model that is designed to acquire categories based purely on stimulus-response associations and error-driven learning can accurately match all of these behavioral profiles, suggesting that an exceptionally simple associative mechanism nicely accounts for the pigeon's manner of learning.

To further examine the learning dynamics that occurred in Phase 2, we analyzed the pigeons' speed of learning as a function of different regions of the stimulus space. The overlap of the stimulus distributions from Phase 1 to Phase 2 presents a unique opportunity for assessing the degree of associative learning expressed by pigeons. [Figure 4A](#) illustrates the stimulus distributions for Phase 1 (Category A in red, Category B in blue) and Phase 2 (Category A in light pink, Category B in light blue); overlaying these distributions are three alpha-transparent squares that divide the stimulus distribution from Phase 2 into three important regions: a transfer region (purple), an interference region (green), and a new region (yellow). In the transfer region, we expect the learned associations from Phase 1 to carry over to Phase 2, such that accuracy should be near the same levels attained at the end of Phase 1. In the new region, because pigeons did not learn any associations in this region during Phase 1, we expect that accuracy should start near chance and gradually increase at a speed similar to learning in Phase 1 (see [Figure 3](#)). In the interference region, the associations from Phase 1 are opposite to those to be learned in Phase 2; therefore, we expect




Figure 3. Data and model fits in the RB/II transfer task

(A) Accuracy data and model fits are shown for the RB (purple) and II (green) tasks across trial blocks in Phase 1. (B) Data and model fits are shown for the RB task across trial blocks in Phase 1 (purple) and Phase 2 (blue). (C) Data and model fits are shown for the RB (blue) and II (yellow) tasks across trial blocks in Phase 2. (D) Data and model fits are shown across trial blocks for the II task in Phase 1 (green) and Phase 2 (yellow). In all panels, learning curves were constructed from daily training sessions by sorting individual pigeon's scores into 10 equally sized blocks of trials and averaging across birds; the resulting data are shown as the solid-colored line with a correspondingly colored shaded region to reflect the standard error of the mean. The model fits are shown as a black line with dots colored to correspond to each condition. RB: rule-based task, II: information integration task, P1: Phase 1, P2: Phase 2.

accuracy to begin well below chance and to rise gradually as the remapping takes place. Perhaps most importantly, if pigeons are using a purely associative learning mechanism, then the patterns of relearning across the three regions should be the same whether or not the task condition is RB or II. We can contrast this prediction with an animal that uses a rule-based mechanism: in the RB task, once the animal begins applying the new rule in Phase 2, accuracy in all three regions ought to rise at the same high rate. However, because there is no single-dimensional rule to learn in the II task, a rule-based mechanism cannot be used; thus, we should see the same performance profile previously described for associative learning.

Figures 4B and 4D depict pigeons' accuracy in Phase 2 separated according to the regions shown in Figures 4A and 4C, where Figures 4A and 4B pertain to the RB task and Figures 4C and 4D pertain to the II task. In Figures 4B and 4D, data are shown as the solid lines colored according to the regions defined in Figures 4A and 4C, respectively, with correspondingly colored shaded regions reflecting the standard error of the mean. Model fits are shown as the black line with dots color coded to correspond with each region. Figures 4B and 4D each depict the profile of associative learning: accuracy starts high and stays high in the transfer condition, accuracy starts around chance and gradually rises in the new region, and accuracy starts well below chance and gradually rises in the interference region. Most importantly, there are no appreciable performance differences between the RB and II tasks – the hallmark of a purely associative learner exhibiting no explicit rule use

Hemispheric RB/II task

The second task explored a similar RB/II design (reported in O'Donoghue et al.²⁶); however, it entailed two important differences in the programmed stimulus distributions, as illustrated in Figures 5A, 5C, 5E, and 5G. First, individual pigeons concurrently learned both RB tasks (involving orientation and frequency rules shown in Figures 5A and 5C, respectively) and both II tasks (involving negative and positive diagonal rules shown in Figures 5E and 5G, respectively); all four tasks were randomly interleaved across trials, with the prevailing task cued by one of the four distinctive background colors. Although this design feature does not explicitly test for associative learning, it would be a remarkable feat if pigeons could accurately use the background colors to modulate the learned associations across the four tasks. Second, the stimulus distributions were abutting, hemispheric distributions rather than the spatially separated elliptical distributions programmed in the prior RB/II task. These abutting distributions allowed us to probe in another way the types of representations pigeons may have used when making their categorization decisions. Here, we systematically examined pigeons' categorization accuracy as a function of the distance the presented training exemplars were from the category boundary. Using an explicit rule that was perfectly aligned with the category boundary ought




Figure 4. Data and model fits by regions of stimulus space

(A) Based on the overlap of the stimulus distributions between Phases 1 and 2, we defined three regions of stimulus space that target different types of learning in Phase 2: a transfer region (purple), an interference region (green), and a new region (yellow). In this panel, the stimulus distributions are shown for the RB task. Red and blue dots correspond to Categories A and B in Phase 1, respectively, whereas light pink and light blue dots correspond to Categories A and B in Phase 2, respectively.

(B) Accuracy curves across trial blocks in Phase 2, separated according to the type of learning region shown in (A) and aggregated across birds in the RB task.

(C) A similar illustration as in (A), but for the II task.

(D) A similar plot as shown in (B), but for the II task. In panels (B) and (D), mean accuracy across blocks is shown as the colored line and the standard error of the mean is shown as the region with corresponding color. Model fits are shown as the black line with dots colored to match the condition of the fit. RB: rule-based task, II: information integration task.

to support a much steeper rise in accuracy than when such dimensional alignment was impossible; thus, the rise in accuracy as a function of the distance from the boundary was predicted to be more precipitous for the RB tasks than for the II tasks. This prediction can be clearly contrasted with the prediction of a purely associative learner; here, there should be no appreciable differences in the rise in accuracy across the RB and II tasks as a function of the distance from the category boundary.

Figures 5B, 5D, 5F, and 5H show the learned associations predicted by the model at the end of training for each of the four tasks corresponding to the stimulus distributions shown in Figures 5A, 5C, 5E, and 5G, respectively. Across all panels, the brightly colored areas designate regions of strong associations, whereas the white areas designate regions of weaker associations. For each of the four tasks, the model accurately learns to associate the regions of the stimulus space with the correct category response. Importantly, the model predicts that the category representations are *graded* in the regions approaching the categorization boundary for each of the four tasks, confirming the dynamic discussed above: both RB and II category boundaries forge associations that support a graded response rule.

Figure 6 shows the data and model fits, aggregated across both of the RB tasks and both of the II tasks. As above, to perform the aggregation, we divided the daily time series associated with each task (frequency RB, orientation RB, negative II, and positive II) into 10 equally spaced bins for each pigeon, calculated the accuracy for those trials, and then averaged over birds and task types. Figure 6A depicts pigeons' accuracy for the RB (purple) and II (green) tasks as a function of trial block, where the data are shown as colored lines and a shaded region to reflect the standard error of the mean, whereas the model fits are shown as the black lines with dots color coded to match the relevant comparison. Over time, we see that pigeons learned the RB and II tasks at very similar speeds, suggesting a purely associative learning mechanism.

Figure 6B shows accuracy as a function of the distance from the category boundary for both the RB (purple) and II (green) tasks. To conduct this analysis, we created 10 equally spaced bins to represent the absolute distance of training stimuli from the categorization boundary for each pigeon, calculated choice accuracy corresponding to those trials (averaging over all time points), and then averaged over pigeons and task types. As with Figure 6A, data are shown as the colored line and shaded region, whereas model fits are shown as the black lines with dots of corresponding color. Figure 6B shows that pigeons' accuracy rises at a very similar rate as a function of the distance from the boundary for both tasks, further indicating that pigeons acquired the RB and II tasks with a purely associative learning mechanism.




Figure 5. Learned representations in the hemispheric task

(A, C, E, G) Stimulus distributions for Category A (red) and Category B (blue) in each of the four different contexts (rows). The different contexts were actually different colors, but here the contexts are designated by different levels of brightness.

(B, D, F, H) Learned associations between the stimulus space and Category A (left panel) and Category B (right panel) in each of the four different contexts defined in panels A, C, E, and G. Brightly colored areas designate strong associations, whereas white areas designate weak associations. Each black hemisphere designates the region in which stimuli from the relevant category occur.

Concentric-rings task

The third task we studied involved another type of categorization structure that requires information integration: the concentric-rings task as reported in O'Donoghue et al.²⁹ In this task, categories are defined by two concentric rings, illustrated in Figures 7A and 7C. For a rule-based learner to acquire this task, they would not only need to use both dimensions, but they would also need to apply a bimodal decision rule along those two dimensions. This type of bimodal response criterion is especially hard to learn relative to unimodal response criteria, especially for humans.^{28–30} However, because a purely associative learner would not need to use such a complex decision rule, they should more readily acquire the task.

Figures 7B and 7D illustrate the associations learned by the model over time. First, Figure 7B shows the associations learned during the first four trials, where two stimuli from each category are depicted (shown in Figure 7A). In both panels, the region of the stimulus distribution is outlined in black; there are two associations to the relevant category (shown as the brighter colors) and two associations to the opposite category (shown as white regions). Although stimuli are associated with the opposing category, associative learning with prediction error can be used to decrease the association between stimuli in that region and that category, thereby producing association strengths that move toward zero. As in Figures 2B and 2D, the spread of association clearly extends beyond the specific values of each training stimulus and into the neighboring regions of the stimulus space, thereby making a strong prediction for categorization responses within these neighboring areas, which we will explain and test later in discussion. Figure 7D illustrates the learned associations at the end of training. Here, Figure 7D shows that, despite the unique shape of the concentric-rings category structure, the model learns to associate regions of the stimulus space with the correct category response.




Figure 6. Data and model fits for the hemispheric task

(A) Accuracy curves across trial blocks for the two RB tasks (purple) and the two II tasks (green).

(B) Accuracy as a function of the distance from the boundary of the two hemispheric distributions for the RB (purple) and II (green) tasks. In both panels, data are shown as the solid-colored line along with the standard error of the mean in the shaded region, whereas model fits are shown as the black lines with dots colored to correspond to each condition. RB: rule-based task, II: information integration task.

Figure 8B depicts the pigeons' accuracy data during training with the depicted stimuli in Figure 8A, aggregated over trial blocks as in the previous two sections. Here, the data are shown as the green solid line, with the shaded region designating the standard error of the mean. The model fits are shown as the black line with green dots. Once again, Figure 8B indicates that the model is well aligned to capture the gradual increase in accuracy over time expressed in the data.

The stimulus distributions in the concentric-rings task were created to eliminate any possible usage of a single, dimensionally aligned decision criterion, making it impossible for rule-based learners to achieve high accuracy. As a result, the distributions afford a compelling advantage to simpler, associative learning mechanisms.

However, to further substantiate the mechanism of associative learning, following training, we used a testing procedure meant to assess the generalization of the categorization responses into the nearby regions created by the ring structure of the training phase (Figure 8A). Figure 8C shows the regions of space corresponding to the conditions of the test phase. First, two conditions simply repeated the stimulus distributions used in the earlier training phase. Second, one form of extrapolation was tested by designing a ring called "OutTest" just beyond the outer ring used in training (light blue region). Third, a form of interpolation was used by defining the region inside the interior training ring called "InTest" (pink region). Fourth, the region between the training regions was defined to examine the gradation of the learned representation, a region simply referred to as Middle (purple region). All five regions were specified such that they were of equal area to control for biases in base rates. Although the InTest and OutTest allowed for clearly correct generalization responses, the Middle region could be associated with both categories learned in the training phase. Therefore, to define the category boundary, we divided the Middle region into two regions of equal area: one associated with the outer ring and one associated with the inner ring.




Figure 7. Learned representations in the concentric-rings task

(A) The first four example stimuli used in the concentric-rings task, along with the outline of the full stimulus distribution; Category A is shown in red and Category B is shown in blue.

(B) Learned associations for Category A (left) and Category B (right) after the first four trials. Brighter colors represent stronger associations, whereas whiter areas indicate association strengths near zero. In both panels, the stimulus distribution corresponding to the relevant category is outlined in black.

(C) The full stimulus distribution used in the concentric-rings task.

(D) Learned representations at the end of training shown in a format similar to (B).




Figure 8. Data and model fits for the concentric-rings task

(A) Stimulus distributions used during training for Category A (red) and Category B (blue).

(B) Accuracy shown as a function of trial blocks. The data are shown as the solid green line with a corresponding shaded region designating the standard error of the mean. The model fits are shown as the black lines with green dots.

(C) Stimulus distributions used during the testing period. Light blue designates the OutTest region, blue designates the OutTrain region (i.e., the region used during training), purple designates the Middle region, red designates the InTrain region (i.e., the region used during training), and pink designates the InTest region.

(D) Accuracy is shown as a function of testing region with the small black lines at the top each bar representing the standard error of the mean. Model fits are shown as the white dots.

Figure 8D depicts the average accuracy of pigeons in the testing phase, separated by the regions of space defined in Figure 8C. Figure 8D indicates that both extrapolation into the OutTest region and interpolation into the InTest region resulted in even *higher* accuracy levels relative to the two training regions – a well-known and unique associative learning effect called “peak shift.”^{31,32} The Middle region serves as a strong baseline, where categorization decisions were only about 60% correct, indicating considerable ambiguity in this region. Across all conditions, the model fits are shown as the black and white dots. Again, a purely associative model accounts extremely well for the profile of accuracy across all five testing regions.

Sectioned-rings task

The final task we studied used stimuli similar to the concentric-rings task, but which subdivided each ring into four segments, as reported in Wasserman et al.²⁷ Figures 9A and 9C depict the stimulus distributions used in the “+ cut” and “x cut” conditions, respectively, where Category A is shown in red and Category B is shown in blue. Like the concentric-rings task, the sectioned-rings task thwarts any effort of the learner to use a simple, unidimensional rule; but the sectioned-rings task eliminates the possibility of a complex rule that would allow a learner to divide the stimulus space into an interior category and an “other” category. Rule formation in this sectioned-rings task would be nearly impossible, yet a pure associative learner ought to acquire the task if given sufficient training.

Figure 9B illustrates the learned representation of the associations formed between the stimulus space and Category A (left) and Category B (right) at the end of training for the + cut condition. Across panels, brighter colors are associated with stronger associations, whereas whiter colors are regions associated with weaker associations. In both panels, the segments that are relevant to each category are outlined in black. In a format similar to Figures 9A–9D illustrate the stimulus distribution and the learned representations at the end of training for the x cut condition. Across both Figures 9B and 9D, the model learns to associate regions of the stimulus space with the correct category response in the outlined region.

Figure 10B depicts the accuracy data and the model fits as a function of trial block. Although we examined the + and x cut conditions separately, because they produced very similar learning profiles, we present the accuracy data aggregated across both tasks. Here, the green line corresponds to the mean accuracy data and the green shaded region corresponds to the standard error of the mean. The model fits are shown as the black line with green dots.




Figure 9. Learned representations in the sectioned-rings task

(A) Stimulus distributions for the “+ cut” condition, where Category A is shown in red and Category B is shown in blue.

(B) Learned associations at the end of training Category A (left) and Category B (right). Brighter colors represent stronger associations, whereas whiter areas indicate near-zero association strength. In both panels, the stimulus distribution corresponding to the relevant category is outlined in black.

(C) Stimulus distributions for the “x cut” condition, where Category A is shown in red and Category B is shown in blue.

(D) Learned representations at the end of training is shown in a format similar to (B).

Figure 10C illustrates the sectioned-rings task shown in Figure 10A, but recoded to divide the space into regions that are near the category boundaries (green) or in the center of the ring sections (purple). This designation allows us to examine the graded nature of the learned representation similar to previous analyses in the hemispheric and concentric-rings tasks. Figure 10D depicts pigeons’ accuracy as a function of trial blocks as shown in Figure 10B, but now separated into the stimulus regions defined by Figure 10C. As in Figure 10B, the data are shown as colored lines with shaded regions designating the standard error of the mean, whereas the model fits are shown as the black lines with correspondingly colored dots. Figure 10D illustrates that both the pigeons and the model learned the middle conditions much more easily than the boundary conditions, once again suggesting a graded associative learning mechanism. As with the earlier tasks, the model fits are in close alignment with the data, providing further evidence that a model involving only associative learning and error minimization satisfactorily accounts for the pigeon’s category learning behavior.

DISCUSSION

In this article, we have clearly documented that pigeons can adeptly acquire a wide variety of category structures, ranging from simple to complex. Indeed, the category structures arranged in the Concentric-Rings Task and the Sectioned-Rings Task were expressly crafted to pose extreme challenges to the deployment of advanced cognitive processes, including selective attention and explicit rule use.

Yet, our pigeons’ ostensible flexibility in solving these highly diverse and demanding tasks may be more apparent than real. The pigeons’ demonstrated success seems to be based on their deploying a single, simple biological mechanism – one that can be effectively emulated by a computational model involving just two free parameters and one hidden layer. That biological mechanism is *associative learning*, which gradually connects behavioral responses to circumscribed regions of perceptual space via an error-correction process. In this sense, the pigeon’s category learning prowess can be understood as if the pigeon were a machine.

One might ask if our “mechanization” of the pigeon is merely a fad in the present era of artificial intelligence. Here, readers might be surprised to learn that efforts to understand animal behavior and cognition by analogy to machines began most famously with the writings of the illustrious French philosopher René Descartes (1596–1650).

In his *Letter to the Marquess of Newcastle*³³ (1646), Descartes argued that animals are nothing more than beastly mechanisms – *bête-machines*.^{34,35} Descartes was willing to concede that animals have sensations and passions; however, he deemed them to be nothing more than organic reactions. Beyond that limited concession, however, Descartes contended that animals lack thought and language; that animals lack abstraction, mathematics, and metacognition³⁶; and, critically, that animals have neither a mind nor a “rational soul,” thus preventing them from engaging in complex cognitive processes.³⁷ For Descartes, the capacity to reason – what he called the *universal instrument*³⁸ – allows humans alone to respond adaptively and flexibly to any and all conditions as well as to achieve mastery over nature.³⁹

Our project puts such contentious philosophical musings concerning mind and reason on the proverbial back burner. To us, treating the pigeon as a machine – *le pigeon machine* – is purely a heuristic enterprise. By doing so, we hope to gain greater insights into the possible biological mechanisms that actually underlie the pigeon’s prolific categorization performance.

Perhaps the most surprising realization of all is that today’s work in machine learning has come to deploy associative algorithms to solve many challenging categorization tasks. Perhaps that’s no accident. Nature’s evolutionary process in creating associative learning may very well serve as an analogous model for computer scientists’ incorporating this powerful algorithm into their own artificial computing systems.^{40,41}




Figure 10. Data and model fits for the sectioned-rings task

(A) Illustration of the stimulus distribution used for Category A (red) and Category B (blue).

(B) Accuracy shown as a function of trial blocks. The data are shown as the solid green line with a corresponding shaded region designating the standard error of the mean. The model fits are shown as the black line with green dots.

(C) Illustration of the stimulus distribution from (A) divided into sections corresponding to the distance from the boundary: stimuli near a boundary are shown in green, whereas stimuli in the middle of the ring sections are shown in purple.

(D) Accuracy is shown as a function of trial block, separated into the regions shown in (C). The data are shown as the solid-colored lines with a corresponding shaded regions designating the standard error of the mean. The model fits are shown as the black lines with correspondingly colored dots.

Limitations of the study

One limitation of the study is that we only explored one type of associative learning model and did not explore additional mechanisms that might underlie the learning profiles of the pigeons. In terms of additional mechanisms, because the pure associative learning mechanism was able to capture all of the qualitative behavioral patterns across the four tasks, it would be unlikely to find quantitative evidence of more sophisticated mechanisms. In terms of alternative associative models, although we chose not to explore other types of associations, we suspect that many of them would prove similarly effective. It is certainly possible that under different experimental conditions, pigeons might exhibit more sophisticated learning strategies that would require deeper investigations into alternative association systems or mechanisms. Because we cannot rule out this possibility, we acknowledge this limitation and save such investigations for future work.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- [KEY RESOURCES TABLE](#)
- [RESOURCE AVAILABILITY](#)
 - Lead contact
 - Materials availability
 - Data and code availability
- [EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS](#)
- [METHOD DETAILS](#)
 - Apparatus
 - Stimuli
 - Procedure
- [QUANTIFICATION AND STATISTICAL ANALYSIS](#)
 - Computational modeling
 - Construction of the learning curves

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2023.107998>.

ACKNOWLEDGMENTS

This research was supported by a CAREER award from the National Science Foundation (BMT), and grant award P01HD080679 (EAW) from the National Institutes of Health. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the article. The authors thank Ellen O'Donoghue, Matt Broschard, Drew Kain, and John Freeman for their contributions to the preparatory work for this article.

AUTHOR CONTRIBUTIONS

E.A.W. designed the research; E.A.W. collected the data; B.M.T. contributed new reagents/analytic tools; B.M.T. analyzed the data; B.M.T. and E.A.W. drafted the article, and both authors contributed significantly to editing the article.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 15, 2023

Revised: June 28, 2023

Accepted: September 18, 2023

Published: September 22, 2023

REFERENCES

- Whitman, C.O. (1919). The Behavior of Pigeons, 3 (Carnegie Institution of Washington. Publication 257).
- Skinner, B.F. (1960). Pigeons in a pelican. *Am. Psychol.* 15, 28–37.
- Capshew, J.H. (1993). Engineering behavior: Project pigeon, World War II, and the conditioning of B. F. Skinner. *Technol. Cult.* 34, 835–857.
- Wasserman, E.A., Brooks, D.I., and McMurray, B. (2015). Pigeons acquire multiple categories in parallel via associative learning: A parallel to human word learning? *Cognition* 136, 99–122.
- Levenson, R.M., Krupinski, E.A., Navarro, V.M., and Wasserman, E.A. (2015). Pigeons (*Columba livia*) as trainable observers of pathology and radiology breast cancer images. *PLoS One* 10, e0141357.
- Navarro, V.M., Wasserman, E.A., and Slomka, P. (2020). Taking pigeons to heart: Birds proficiently diagnose human cardiac disease. *Learn. Behav.* 48, 9–21.
- Morgan, M.J., Fitch, M.D., Holman, J.G., and Lea, E.G. (1976). Pigeons learn the concept of an "A". *Perception* 5, 57–66.
- Watanabe, S., Lea, S.E.G., and Dittrich, W.H. (1993). What can we learn from experiments on pigeon concept discrimination? In *Vision, Brain, and Behavior in Birds*, 1, H.P. Zeigler and H.-J. Bischof, eds. (MIT Press), pp. 351–376.
- Herrnstein, R.J. (1985). Riddles of natural categorization. *Philosophical Transactions of the Royal Society B308*, 129–144.
- Katz, J.S., Wright, A.A., and Bodily, K.D. (2007). Issues in the comparative cognition of abstract-concept learning. *Comp. Cogn. Behav. Rev.* 2, 79–92.
- Qadri, M.A.J., Ashby, F.G., Smith, J.D., and Cook, R.G. (2019). Testing analogical rule transfer in pigeons (*Columba livia*). *Cognition* 183, 256–268.
- Troje, N.F., Huber, L., Loidolt, M., Aust, U., and Fieder, M. (1999). Categorical learning in pigeons: The role of texture and shape in complex static stimuli. *Vis. Res.* 39, 353–366.
- Wasserman, E.A., Kiedinger, R.E., and Bhatt, R.S. (1988). Conceptual behavior in pigeons: Categories, subcategories, and pseudocategories. *J. Exp. Psychol. Anim. Behav. Process.* 14, 235–246.
- Watanabe, S. (2011). Discrimination of painting style and quality: Pigeons use different strategies for different tasks. *Anim. Cognit.* 14, 797–808.
- Güntürkün, O. (2020). The Surprising Power of the Avian Mind. *Sci. Am.* 322, 48–55.
- Zorina, Z.A., Obozova, T.A., and Smirnova, A.A. (2022). Higher cognitive abilities in birds: A comparative evolutionary analysis. *Neurosci. Behav. Physiol.* 52, 62–76.
- Stacho, M., Herold, C., Rook, N., Wagner, H., Axer, M., Amunts, K., and Güntürkün, O. (2020). A cortex-like canonical circuit in the avian forebrain. *Science* 369, eabc5534.
- Pusch, R., Clark, W., Rose, J., and Güntürkün, O. (2023). Visual categories and concepts in the avian brain. *Anim. Cognit.* 26, 153–173.
- Ashby, F.G., Alfonso-Reese, L.A., Turken, A.U., and Waldron, E.M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* 105, 442–481.
- Lind, J. (2018). What Can Associative Learning Do for Planning? 5 (Royal Society Open Science), 180778.
- Ibarz, J., Tan, J., Finn, C., Kalakrishnan, M., Pastor, P., and Levine, S. (2021). How to train your robot with deep reinforcement learning: Lessons we have learned. *Int. J. Robot Res.* 40, 698–721.
- Wurman, P.R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T.J., Capobianco, R., Devlic, A., Eckert, F., Fuchs, F., et al. (2022). Outracing champion gran turismo drivers with deep reinforcement learning. *Nature* 602, 602.
- Smith, J.D., Berg, M.E., Cook, R.G., Murphy, M.S., Crossley, M.J., Boomer, J., Spiering, B., Beran, M.J., Church, B.A., Ashby, F.G., and Grace, R.C. (2012). Implicit and explicit categorization: A tale of four species. *Neurosci. Biobehav. Rev.* 36, 2355–2369.
- Maddox, W.T., and Ashby, F.G. (2004). Dissociating explicit and procedural-learning based systems of perceptual category learning. *Behav. Process.* 66, 309–332.
- Smith, J.D., Ashby, F.G., Berg, M.E., Murphy, M.S., Spiering, B., Cook, R.G., and Grace, R.C. (2011). Pigeons' categorization may be exclusively nonanalytic. *Psychon. Bull. Rev.* 18, 414–421.
- O'Donoghue, E.M., Broschard, M.B., and Wasserman, E.A. (2020). Pigeons exhibit flexibility but not rule formation in dimensional learning, stimulus generalization, and task switching. *J. Exp. Psychol. Anim. Learn. Cogn.* 46, 107–123.
- Wasserman, E.A., Kain, A.G., and O'Donoghue, E.M. (2023). Resolving the associative learning paradox by category learning in pigeons. *Curr. Biol.* 33, 1112–1116.e2.
- Turner, B.M. (2019). Toward a common representational framework for adaptation. *Psychol. Rev.* 126, 660–692.
- O'Donoghue, E.M., Broschard, M.B., Freeman, J.H., and Wasserman, E.A. (2022). The lords of the rings: People and pigeons take different paths mastering the concentric rings categorization task. *Cognition* 218, 104920.

30. Turner, B.M., Van Zandt, T., and Brown, S. (2011). A dynamic, stimulus-driven model of signal detection. *Psychol. Rev.* 118, 583–613.
31. Purtle, R.B. (1973). Peak shift: A review. *Psychol. Bull.* 80, 408–421.
32. Spence, K.W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychol. Rev.* 44, 430–444.
33. Ariew, R. (2000). Rene Descartes: Philosophical Essays and Correspondence'.
34. Gunderson, K. (1964). Descartes, La Mettrie, language, and machines. *Philosophy* 39, 193–222.
35. Newman, L. (2001). Unmasking Descartes' case for the bête machine doctrine. *Can. J. Philos.* 31, 389–425.
36. Smith, C.U. (2005). Book review essay. How the modern world began: Stephen Gaukroger's Descartes' system of natural philosophy. *J. Hist. Neurosci.* 14, 57–63.
37. Avramides, A. (1996). Descartes and other minds. *Teorema XVI*, 27–46.
38. Wilson, M.D. (1995). Animal ideas. *Proceedings and Addresses of the American Philosophical Association* 69, 7–25.
39. Melehy, H. (2005). Silencing the animals: Montaigne, Descartes, and the hyperbole of language. *Symploke* 13, 263–282.
40. Wen, H., Fan, H., Xie, W., and Pei, J. (2017). Hybrid structure-adaptive RBF-ELM network classifier. *IEEE Access* 5, 16539–16554.
41. Wen, H., Xie, W., and Pei, J. (2016). A structure-adaptive hybrid RBF-BP classifier with an optimized learning strategy. *PLoS One* 11, e0164719.
42. Nelder, J.A., and Mead, R. (1965). A simplex method for function minimization. *Comput. J.* 7, 308–313.
43. Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., and Broussard, C. (2007). What's new in psychtoolbox-3. *Perception* 36, 1–16.
44. Logan, G.D. (1988). Toward an instance theory of automatization. *Psychol. Rev.* 95, 492–527.
45. Cohen, J.D., Dunbar, K., and McClelland, J.L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychol. Rev.* 97, 332–361.
46. Gluck, M.A., and Bower, G.H. (1988). Evaluating an adaptive network model of human learning. *J. Mem. Lang.* 27, 166–195.
47. Pearce, J.M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychol. Rev.* 101, 587–607.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Rule-based and Information Integration Switch Tasks	O'Donoghue et al. ²³	https://osf.io/3265r/
Hemispheric Distributions Task	O'Donoghue et al. ²³	https://osf.io/3265r/
Concentric Rings Task	O'Donoghue et al. ²⁶	https://osf.io/vyd3p/
Sectioned Rings Task	Wasserman et al. ³⁸	https://osf.io/b639r/
Software and algorithms		
Optimization function in R	Nelder and Mead ⁴²	https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/optm
Other		
Code used for fitting the model to data	This paper	https://osf.io/b4eac/

RESOURCE AVAILABILITY

Lead contact

Further information or requests for materials should be directed to and will be fulfilled by the lead contact, Brandon Turner (turner.826@gmail.com).

Materials availability

No materials were generated for this study.

Data and code availability

- This paper analyzes existing, publicly available data. These accession numbers for the datasets are listed in the [key resources table](#).
- All original code has been deposited at <https://osf.io/b4eac/> and is publicly available as of the date of publication. DOIs are listed in the [key resources table](#).
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

We trained a total of 24 adult male and female pigeons (*Columba livia*) on the four different categorization tasks: 8 on the RB/II Transfer Task, 4 on the Hemispheric RB/II Task, 8 on the Concentric-Rings Task, and 4 on the Sectioned-Rings Task. All of the pigeons had varied prior experience in our laboratory, but they had no experience with the present tasks or visual stimuli. Following the procedures approved by the Institutional Animal Care and Use Committee at The University of Iowa, the pigeons were individually housed and maintained at 85% of their free-feeding weights.

METHOD DETAILS

Apparatus

The experiments were conducted in 36 × 36 × 41 cm operant conditioning chambers with white noise sounded during all sessions. The stimuli were displayed on 800 × 600 px LCD monitors, and the pigeons' pecking responses were transduced by AccuTouch touchscreens (Elo TouchSystems, Fremont, CA). The screen area visible to the pigeons measured 28.5 × 17.0 cm. Correct responses were reinforced with one to three 45-mg pigeon pellets delivered via a rotary dispenser connected to a food tray mounted opposite the touchscreen and monitor.

Stimuli

The visual stimuli were black-and-white circular sinusoidal gratings that varied in spatial frequency and line orientation. The stimuli were 6.29 × 6.29 cm in the RB/II Transfer Task, 2.64 × 2.64 cm in the Hemispheric RB/II Task, 7.50 × 7.50 cm in the Concentric-Rings Task, and 7.50 × 7.50 cm in the Concentric-Rings Task. All stimuli were randomly generated using Matlab® and its Psychtoolbox-3 extensions.⁴³

In each training session, we uniformly and randomly sampled stimuli from the relevant sections of the bidimensional stimulus space: 80 stimuli from Category A and 80 stimuli from Category B. Testing sessions were modified so that a sufficient number of stimuli from the training and testing regions could be available for statistical purposes.

Procedure

At the start of each daily trial (illustrated in Figure S1), a white star-shaped ‘start stimulus’ overlaid on a black square was presented in the center of the screen, surrounded by a solid gray background. Once the pigeon pecked the start stimulus, it disappeared, and a randomly sampled ‘trial exemplar’ (i.e., stimulus relevant to the task) appeared in its place. The pigeons had to complete a number of observing pecks to this trial exemplar, beginning with 1 peck and adjusted upward across sessions according to each pigeon’s performance to a maximum of 20 pecks. Once the observing response requirement was met, the trial exemplar disappeared and two replicas of the trial exemplar appeared on opposite sides of the screen from where the trial exemplar had just been presented. These replicas served as the choice buttons, with the correct and incorrect buttons for each category randomly located on opposite sides of the screen from where the trial exemplar had just been presented. If the pigeon pecked the correct choice button, then that button was replaced by a white square; a single peck to the white square yielded food reinforcement. If the pigeon pecked the incorrect choice button, then a correction trial was initiated following a variable 6- to 10-s delay. Correction trials continued to be given until the pigeon made the correct response, ensuring that all trials ended with food reinforcement. Correction trials were excluded from data analyses. A variable 6- to 10-s period served as the intertrial interval.

Each pigeon was trained until it reached a criterion level or its own performance asymptote had been attained. Learning curves were constructed from these training sessions by sorting scores into 10 equally-sized blocks of trials.

QUANTIFICATION AND STATISTICAL ANALYSIS

Computational modeling

The model we developed in the present manuscript can be considered a multivariate extension of the adaptive representation model (ARM) detailed in Turner.²⁸ There, a set of learning models was developed with different assumptions about learning mechanisms to relate instance learning models^{30,44} to strength-based models commonly used in connectionist architectures.^{45–47} One model in the set that was particularly successful in explaining relearning dynamics was the so-called strength-prediction error (SEP) model that combined simple associative learning with prediction error to associate regions of the stimulus space to categorization decisions. Because this model possessed the two mechanisms (i.e., stimulus-response association and stimulus generalization) hypothesized to underlie pigeon learning behavior, we adapted it for our purposes here.

The model contains a grid of “representational points” over the stimulus space, in our case, over the bidimensional space of possible frequency and orientation values. Let X denote the matrix of grid values, where each row contains a location in the stimulus space. For our purposes, we assumed a grid of points ranging from 0 to 100 in increments of 20 along both dimensions, meaning that there were $n = 400$ representation points and that X was a (400×2) matrix. As illustrated in Figure 1, associated with each representation point in X is an associative strength to each category, which changes from trial to trial. Let P_t denote the matrix of associative strengths on Trial t . In our case, each task consisted of only two category responses, so each P_t was a (400×2) matrix. For the hemispheric task, there were four different contexts, and so we expanded P_t to be a $(400 \times 2 \times 4)$ array, and then separately updated the appropriate matrix based on the context experienced on a given trial. For notational convenience, we will ignore this multidimensional contextual array case in the equations that follow.

We first constructed the mapping between the stimulus presented on Trial t and the hidden layer containing the representation points in X . To do this, we assumed a similarity kernel where regions of the stimulus space were related to the representation points via

$$K(x_i | s_t) = e^{-\delta \sum_k |x_{i,k} - s_{t,k}|}$$

where $x_{i,k}$ denotes the k th dimension of the i th representation point ($i \in \{1, 2, \dots, n\}$), and $s_{t,k}$ denotes the value of the stimulus presented on Trial t on the k th dimension. Evaluating this equation produces a value that represents how close the stimulus s_t is to the i th representation point x_i . Figure 1B illustrates how this equation relates to the stimulus value.

The value of the similarity kernel determines how relevant the i th representation point is when deciding about the stimulus s_t . However, to determine the evidence for each category, we combined relevance with the set of learned associations up to Trial t . To do this, we simply multiplied the relevance value by the association strength for each category to obtain the evidence for Category k :

$$E_t(k) = \sum_i K(x_i | s_t) P_{t,i,k}$$

where $P_{t,i,k}$ denotes the association of the i th representation point to Category k on Trial t . We assumed that the probability of responding Category k on Trial t , denoted $Pr_t(k)$, was calculated by evaluating the ratio of the evidence for Category k to the sum of evidence for all categories via

$$Pr_t(k) = \frac{E_t(k)}{\sum_z E_t(z)}$$

Following the decision, feedback was given to the pigeon (and the model). Letting $f_{t,k}$ denote the feedback on Trial t , we arbitrarily assigned Category A feedback to.

Dimension 1 (i.e., $k = 1$) and Category B feedback to Dimension 2 (i.e., $k = 2$). The value of $f_{t,k}$ then corresponded to which category was reinforced such that

$$f_t = \begin{cases} [1 & 0] & \text{if Category A reinforced} \\ [0 & 1] & \text{if Category B reinforced} \end{cases}$$

This vector of feedback is then used as the target function in the Rescorla-Wagner updating rule over the set of learned associations. We updated the associations via the following learning rule:

$$P_{t+1,i,k} = P_{t,i,k} + \alpha(f_{t,k} - P_{t,i,k})K(x_i|s_t)$$

Based on this specification, the model contains only two parameters: the learning rate parameter α and the stimulus generalization parameter δ . We separately fit the model to each pigeon's data by using the `optim` function in R to obtain the maximum likelihood estimate for both parameters.⁴² The likelihood function was specified by relating the model predicted probabilities of each response to the observed pigeon response. Letting R_t denote the pigeon's response on Trial t such that (arbitrarily)

$$R_t = \begin{cases} 0 & \text{if Category A response} \\ 1 & \text{if Category B response} \end{cases}$$

then the likelihood function was defined as

$$L(\alpha, \delta|R) = \prod_t \text{Bernoulli}(R_t|P_{r_t}(B))$$

where $\text{Bernoulli}(a|b)$ denotes the Bernoulli density function for outcome a with probability b .

Construction of the learning curves

Because the number of days to meet learning criterion differed for different pigeons, we partitioned each bird's daily scores into the same number of scoring blocks so that all birds would contribute equally to the fitted functions. This process did not affect the overall conclusions we reached. For example, in the first phase of the RB/II Transfer Task, all eight pigeons reached the 85% criterion between 3 and 9 sessions. The mean number of sessions in the RB group was 4.25, SD = 1.50 and the mean number of sessions in the II group was 4.50, SD = 1.00. The number of sessions required to reach criterion did not reliably differ between groups, $t(3.17) = 0.82$, $p = 0.468$, indicating that the nature of the two categorization tasks did not differentially affect learning speed. In the second phase of the RB/II Transfer Task, all eight pigeons reached the 85% criterion within 6 sessions; the mean number of sessions required to reach criterion was 4.25 for both the RB (SD = 0.96) and II (SD = 1.50) groups. Sessions to criterion did not differ as a function of task, $t(5.10) = 0.00$, $p = 1.000$. Once again, these data indicate that the nature of the two categorization tasks did not differentially affect learning speed.