Published in final edited form as:

Psychol Sci. 2010 October 1; 21(10): 1510–1517. doi:10.1177/0956797610382125.

Missing the Forest for the Trees: Object-Discrimination Learning Blocks Categorization Learning

Fabian A. Soto and Edward A. Wasserman

University of Iowa

Abstract

Growing evidence indicates that error-driven associative learning underlies the ability of nonhuman animals to categorize natural images. This study explored whether this form of learning might also be at play when people categorize natural objects in photographs. Two groups of college students (a blocking group and a control group) were trained on a categorization task and then tested with novel photographs from each category; however, only the blocking group received pretraining on a task that required the discrimination of objects from the same category. Because of this earlier noncategorical discrimination learning, the blocking group performed well in the categorization task from the outset, and this strong initial performance reduced the likelihood of category learning driven by error. There was far less transfer of categorical responding during testing in the blocking group than in the control group; this finding suggests that learning the specific properties of each photographic image in pretraining blocked later learning of an open-ended category.

Keywords

object categorization; natural scenes; associative learning; Rescorla-Wagner theory

In order to adapt to their environment, organisms must learn to respond similarly to nonidentical objects from the same category (categorical responding) as well as to respond differently to similar objects from the same category (noncategorical discrimination). Underlying such learning must be a psychological mechanism that detects and extracts those aspects of objects that are invariant across class members, to support categorization, as well as those aspects that are specific to each stimulus, to support discrimination. Once these different stimulus properties are extracted, they should control behavior in a task depending on whether they are more useful or less useful for achieving adaptive performance.

An important aspect of categorization research is understanding how animals learn to classify photographs of natural objects (Lazareva & Wasserman, 2008). We recently proposed a model of this categorization behavior based on principles derived from associative-learning theory (Soto & Wasserman, 2010). The idea behind this model is simple: Animals represent each image as a collection of “elements,” which vary in their level of specificity and invariance with respect to the stimuli they embody. Some elements tend to be activated by a single image only—these elements represent the image’s stimulus-specific properties; other elements are activated by several different images depicting objects from the same category—these elements represent category-specific properties of the image. In a categorization task, these stimulus properties are associated with responses depending on the ability of those properties to predict reward via an error-driven learning

rule (Rescorla & Wagner, 1972). Stimulus-specific properties control performance in discrimination tasks, whereas category-specific properties control performance in categorization tasks.

Error-driven learning is especially useful in situations in which organisms must predict important outcomes by assessing signaling stimuli. In a categorization task, animals must predict which response to a particular stimulus will produce food reward. How much learning occurs on each trial is proportional to the degree of error made in predicting the outcome. If the outcome is unexpected, then the preceding stimulus should acquire a strong association with that outcome; but if the outcome is perfectly predicted, then no learning about the stimulus should occur.

Researchers generally agree that error-driven learning plays an important part in simple forms of associative learning, such as Pavlovian and instrumental conditioning (Schwartz, Wasserman, & Robbins, 2002). Many different models of associative learning (reviewed by Vogel, Castro, & Saavedra, 2004) are based on an error-driven learning rule. The same is true of models of more complex cognitive phenomena, including connectionist networks (Rumelhart & McClelland, 1986) and reinforcement-learning models (Frank & Daw, 2009). In addition, there is growing neurobiological evidence that prediction errors are computed in the brain and used for learning and other adaptive functions (Maia, 2009; Niv, 2009; Schultz & Dickinson, 2000).

Error-driven learning explains several results in the literature on natural-image categorization by animals; it also generates new predictions about the conditions that should foster or hinder categorization learning (see Soto & Wasserman, 2010). One key prediction is that animals may be impaired in learning that a set of stimuli belongs to a category if they have learned previously to discriminate among those stimuli.

Table 1 outlines an experimental design devised to test this prediction. The experimental condition is termed *blocking*. Participants are trained in Phase 1 on a pseudocategorization task, in which photographs of objects from two basic-level categories are arbitrarily assigned to one of two different responses. This task cannot be solved by relying on perceptual resemblance among stimuli in the same class; participants must learn instead to discriminate objects from the same basic-level category. In Phase 2, training continues with only half of the previously trained stimuli, but now in a true categorization task, in which all of the photographs from the same category are assigned to the same response. Because this stimulus-response mapping was trained earlier in the pseudocategorization task, participants should maintain good discrimination behavior and make few, if any, errors in predicting the correct response for each of the stimuli. If learning is driven by prediction error, then learning the mapping of categories to responses should be impaired. In contrast, the type of training in Phase 2 should lead to robust categorization learning if pseudocategorization training does not precede it, as in the control condition for Experiment 1, detailed in Table 1.

Figure 1a shows the predictions of our error-driven learning model for this experiment. (A more detailed description of this model and the parameter values used in the simulation can be found in Soto & Wasserman, 2010.) The model predicts that, during testing with both stimuli previously used in training and novel stimuli, the control condition should foster high performance with the training stimuli, as well as robust transfer of learning to novel exemplars from each category. Such transfer is usually taken as evidence of categorization learning (Bhatt, Wasserman, Reynolds, & Knauss, 1988; Herrnstein, 1990). In contrast, the blocking condition should produce much lower transfer of learning to novel stimuli from the trained categories. As shown in Figure 1b, this prediction proved to be true for pigeons (Soto & Wasserman, 2010). The qualitative pattern of results was the same in the observed data

and the simulated data: Figures 1a and 1b both show a disparity in transfer behavior that is analogous to the blocking effect in classical conditioning (Kamin, 1969)—this effect is generally considered evidence of error-driven learning.

As for humans, research in object recognition has largely focused on how the visual system extracts shape information from images and builds object representations. Theories in this area place little emphasis on the role of supervised learning and decision processes in categorization (Palmeri & Gauthier, 2004). Many contemporary models of object recognition (Hummel, 2001; Hummel & Stankiewicz, 1998; Riesenhuber & Poggio, 1999; Serre, Oliva, & Poggio, 2007) do not include error-driven learning among their key mechanisms. Thus, it is an open question whether error-driven learning participates in human categorization of natural objects, as seems to be the case with other animals.

Research with artificial categories has provided some direct evidence for error-driven learning in visual categorization (Gluck & Bower, 1988; Nosofsky, Kruschke, & McKinley, 1992; Shanks, 1991). But the main focus in this area of research has been the structure of category-knowledge representations and how such representations are used in decision making (Palmeri & Gauthier, 2004). In this area, too, many popular theories do not incorporate error-driven learning (e.g., Ashby, 1992; Nosofsky, 1984), and assumptions about category learning are only rarely explored. It is also unclear how the learning principles that have been proposed to explain artificial categorization might participate in the categorization of natural objects.

The prime aim of the present study was to determine whether category learning is blocked in humans, as it is in pigeons, by prior noncategorical-discrimination learning in a task involving photographs of natural objects. Evidence of this effect would suggest that the same learning principles exhibited in animals also participate in object categorization in people.

To make the results of this study comparable with those of prior animal research, we kept the stimuli and procedures as similar as possible to the methods used with pigeons. Using the same stimuli as in previous animal research (photographs of cars, chairs, flowers, and people; see Fig. 2) was also interesting because it seems natural for people to represent such stimuli very differently than pigeons do. For pigeons, these photographic stimuli are entirely novel. For people, these stimuli are generally familiar; each of these images has a readily available verbal label, which allows learning of a simple classification rule of the form “if the photograph depicts a car, then respond ‘left.’” Despite any differences in the way humans or pigeons represent these stimuli, the experimental design used in our study was created to test whether the principles of category learning are the same for both species, that is, whether the absence of prediction error impairs categorization learning in people as it does in pigeons.

Experiment 1

This experiment was designed to test the blocking of category learning in people by using the same design previously used with pigeons.

Method

Participants—Forty-eight undergraduates (males and females) from the University of Iowa participated in exchange for course credit. They had normal or corrected-to-normal vision.

Stimuli and apparatus—The stimuli were 120 color photographs depicting 30 objects from four categories (people, flowers, cars, and chairs; see Fig. 2 for black-and-white examples) against varied backgrounds. Only photographs from two categories were shown to each participant. Across participants, all four of the categories were used equally often in all possible two-category combinations, and the assignment of categories to response keys was counterbalanced.

The experiment was conducted on five Apple eMac computers. The procedure was programmed using Matlab 7.0.4 (MathWorks, Natick, MA). Images of the photographs were displayed on a 107.0-cm × 70.5-cm area in the middle of the computer monitor. The response keys were two 2.5-cm × 2.5-cm icons displaying different black-and-white patterns. They were positioned on the screen immediately to the left and to the right of the display area. Auditory feedback was provided to participants through headphones.

Procedure—Participants were given written instructions that they would be presented with several photographs and asked to choose which of the two response keys was associated with each picture. Participants were told that they would at first have to guess which response was correct, but that they would receive feedback after every trial to help them improve their accuracy. Participants were asked to respond as quickly and accurately as possible.

Each trial began with a 1-s presentation of a white cross in the middle of the screen. Immediately afterward, a single photograph was presented on the screen. After 200 ms, response keys appeared—one to the left and one to the right of the photograph—and the participant had to mouse-click on one of the keys to advance the trial. Feedback about the participant's choice was provided to him or her for 1 s at the end of each trial: "Correct" or "incorrect" was displayed on the screen, and a pleasant chime or an unpleasant buzzer, respectively, was sounded via the headphones. There was a 1-s interstimulus interval before the next trial.

The experimental design included the blocking and control conditions shown in Table 1. Half of the participants were randomly assigned to each condition. In each phase of each condition, a training block involved one presentation of each trial type shown in Table 1. Each trial type was repeated 10 times during a block; each time a different photograph from a single category was presented, and a correct response was assigned exclusively to either the left or the right response key. Trials were randomized within blocks.

Participants in the blocking group were given 10 blocks of training in Phase 1, immediately followed by five blocks of training in Phase 2. The testing phase started immediately after the end of Phase 2; it involved 10 trials presenting different novel stimuli from each of the two training categories intermixed with all trials of one block of Phase 2 training. Participants in the control group were given the same training and testing as participants in the blocking condition, but without Phase 1 training.

Results and discussion

Figure 1c shows the mean proportion of correct choices in the testing phase for the two groups. Performance on Phase 2 training trials was very high and nearly identical in the two groups; however, the generalization of categorization learning differed greatly. The mean proportion of correct choices in response to novel stimuli was high for participants in the control group ($M = .97$, $SE = .03$), but considerably lower for participants in the blocking group ($M = .68$, $SE = .09$).

The testing data were analyzed with a 2 (group: blocking or control) $\times$ 2 (trial type: training stimuli or novel stimuli) $\times$ 12 (counterbalancing) analysis of variance (ANOVA), with choice accuracy as the dependent variable and participant as a random factor. There was a significant interaction between group and trial type, $F(1, 24) = 56.99$, $p < .001$, $\eta_p^2 = .70$. Post hoc tests (Newman-Keuls, $\alpha = .05$) disclosed that the disparity between groups was significant on trials with novel stimuli, $d = 0.80$, but not on trials with training stimuli, $d = 0.06$. Furthermore, the disparity between performance with training stimuli and performance with novel stimuli was significant in the blocking group, $d = 0.92$, but not in the control group, $d = 0.23$.

The ANOVA also revealed a significant main effect of trial type, $F(1, 24) = 80.68$, $p < .001$, $\eta_p^2 = .77$; this effect reflected the fact that both groups showed some decrement in performance on the novel stimuli. Finally, there was a significant main effect of group, $F(1, 24) = 48.08$, $p < .001$, $\eta_p^2 = .67$; this effect was a consequence of the higher overall performance level of the control group. No other effects in the ANOVA were significant.

These results suggest that for humans, just as for pigeons, learning to discriminate objects from the same category impairs learning to sort the same objects into their basic-level categories. Without prior discrimination learning, categorization learning generalizes almost perfectly to novel exemplars. This blocking effect is the first evidence indicating that error-driven associative learning is involved in human categorization of natural objects.

Experiment 2

Although the results of Experiment 1 were the direct prediction of a model based on error-driven learning, they can be explained in another way. Stimuli from the same category were paired with more than one response only in the blocking condition, whereas in the control condition, stimuli from the same category were always assigned to a single response. The results may have been due to a decision process occurring in testing, rather than a learning effect during training. People in the control group might have shown a high level of categorical responding simply because they saw only one mapping of categories to responses in training and generalized that single mapping to the novel testing stimuli. In comparison, people in the blocking group might have shown a lower level of categorical responding because they saw an inconsistent mapping of categories to responses in training.

This alternative explanation would be consistent with a model that stores the frequencies of event co-occurrences and classifies stimuli accordingly, as do some probabilistic models of associative learning (e.g., Cheng, 1997) and some decision-based models of perceptual categorization (e.g., Ashby, 1992; Nosofsky, 1984). These theories are sensitive to conflicting information about the mapping between categories and responses, but they are insensitive to the temporal order in which that information is acquired; they assume that people memorize all of the co-occurrence information and later use it to make a decision, with each piece of information having the same impact over the final choice regardless of when it was received. In contrast, error-driven learning theories are sensitive to the temporal order in which conflicting information is received, and this gives rise to trial-order effects. According to these theories, the most recently presented information has the greatest impact on behavior.

Participants in the control condition of Experiment 2 (Table 1) were trained with one mapping of categories onto responses in Phase 1, but the opposite mapping in Phase 2. An error-driven learning model would predict that when shown new category exemplars in testing, participants should respond to them according to the category-response mapping learned most recently (in this case, in Phase 2). Thus, although the control condition in Experiment 2 involved an inconsistent category-response mapping across training phases, a

model based on error-driven learning would still predict more categorical responding to the novel testing stimuli in this control condition than in the blocking condition.

In contrast, a decision model based on frequencies of co-occurrences would give the same weight to the information in the two phases and predict no preference for one category-response mapping over another. Thus, according to the decision model, the level of categorical responding to the novel testing stimuli should be similarly low in both the control and the blocking conditions. Experiment 2 tested which of these predictions would prevail.

Method

Participants were 72 undergraduates similar to the subjects in Experiment 1. Half were randomly assigned to the blocking group, and the other half to the control group. Stimuli, apparatus, and procedures were as described for Experiment 1, with the exception that participants in the control group were given five blocks of training in Phase 1, and the assignment of stimuli to responses in that group was opposite in Phase 1 and Phase 2.

Results and discussion

The main results from testing are shown in Figure 1d. The pattern of results was the same as observed in Experiment 1: high performance for Phase 2 training trials in both groups and a disparity between groups in the level of transfer to the novel stimuli. The mean proportion of correct choices in response to novel stimuli was high in the control group ($M = .96$, $SE = .03$), but much lower in the blocking group ($M = .67$, $SE = .08$).

The testing data were analyzed with a 2 (group: blocking or control) $\times$ 2 (trial type: training stimuli or novel stimuli) $\times$ 12 (counterbalancing) ANOVA, with choice accuracy as the dependent variable and participant as a random factor. A statistically significant interaction between group and trial type, $F(1, 48) = 58.39$, $p < .001$, $\eta_p^2 = .55$, supported the described pattern of results. Post hoc tests (Newman-Keuls, $\alpha = .05$) indicated that the disparity between groups was significant for novel stimuli, $d = 0.81$, but not for training stimuli, $d = 0.02$. Also, the disparity between trials with training stimuli and trials with novel stimuli was significant in the blocking group, $d = 0.94$, but not in the control group, $d = 0.19$. As in Experiment 1, there were significant effects of trial type, $F(1, 48) = 84.21$, $p < .001$, $\eta_p^2 = .64$, and group, $F(1, 48) = 61.06$, $p < .001$, $\eta_p^2 = .56$. No other effects were significant.

These results replicate the blocking effect found in Experiment 1; they also refute an explanation of that effect based on the inconsistent assignment of stimuli to responses in the blocking group. In this experiment's control group, participants generalized the category-response mapping that they learned in Phase 2 to novel stimuli, and they did this even when they had experienced the opposite mapping in Phase 1. This result is inconsistent with an explanation of the blocking effect by decision models that propose that information is stored in the form of co-occurrence frequencies; instead, this result is consistent with an explanation in terms of error-driven learning.

Experiment 3

In Experiments 1 and 2, we inconsistently assigned categories to responses in the same training phase only in the blocking condition; we did this to encourage memorization of each individual stimulus and its corresponding response. Under such conditions, however, participants might simply learn that whenever new stimuli are presented, they should try to memorize their assigned responses. This approach would lead to random responding to novel stimuli in testing. To eliminate this possibility, we designed a different control condition in Experiment 3 (Table 1): Participants were trained on a pseudocategorization

task in Phase 1, just as in the blocking condition. If what people learn in pseudocategorization training is a general “memorization” strategy that influences later choices, then participants in both the control condition and the blocking condition should show similar test behavior.

However, in this new control condition, Phase 2 training was carried out with novel stimuli from the training categories. Our account of category learning predicts that the novel stimuli shown in Phase 2 should produce prediction error and therefore prompt error-driven category learning that should generalize in testing, just as in the control conditions of Experiments 1 and 2.

Method

Participants were 72 undergraduates similar to the participants in Experiments 1 and 2. Half were randomly assigned to the blocking condition, and half to the control condition. Stimuli, apparatus, and procedures were the same as in Experiment 1, with the exception that participants in the control group were exposed to 10 blocks of training in Phase 1. This training involved a pseudocategorization task with stimuli that were different from the stimuli presented afterward in Phase 2.

Results and discussion

Results from testing are shown in Figure 1e. The pattern of results was the same as in Experiments 1 and 2. Performance on Phase 2 training trials was similarly high in the two groups. Testing performance on novel stimuli was high in the control group ($M = .91$, $SE = .05$), but much lower in the blocking group ($M = .57$, $SE = .08$).

The testing data were analyzed with a 2 (group: blocking or control) $\times$ 2 (trial type: training stimuli or novel stimuli) $\times$ 12 (counterbalancing) ANOVA, with choice accuracy as the dependent variable and participant as a random factor. There was a significant interaction between group and trial type, $F(1, 48) = 60.31$, $p < .001$, $\eta_p^2 = .56$. Post hoc tests (Newman-Keuls, $\alpha = .05$) indicated that the groups differed significantly on trials with novel stimuli, $d = 0.85$, but not on trials with training stimuli, $d = 0.03$. Performance differed significantly between trials with novel stimuli and trials with training stimuli for the blocking group, $d = 1.14$, but not for the control group, $d = 0.29$. There were also significant effects of trial type, $F(1, 48) = 114.44$, $p < .001$, $\eta_p^2 = .71$, and group, $F(1, 48) = 39.58$, $p < .001$, $\eta_p^2 = .45$. No other effects were significant.

These results replicate the blocking effect found in the two prior experiments and eliminate the possibility that mere pretraining on a pseudocategorization task produces this effect. As predicted by error-driven learning, the blocking effect appears only if training in the categorization task involves the same stimuli previously experienced in a pseudocategorization task. That is, the effect is stimulus-specific and cannot be explained by learning a general memorization strategy.

General Discussion

In three experiments, we found that participants' categorization learning was impaired if they learned to discriminate objects from the same category prior to being instructed to sort those objects by that category; this observation was revealed by a generalization test with novel category exemplars. In Experiment 1, the results of the control condition showed that the same categorization training the blocking group received, but without the prior discrimination training given to the blocking group, supported robust generalization to novel exemplars. In other words, when people first had to memorize individual objects and their arbitrarily assigned responses, they subsequently found it difficult to detect a change when

all of the presented objects were sorted according to their basic-level categories; to put it colloquially, they “missed the forest for the trees.”

The results of Experiments 2 and 3 further suggest that the blocking effect is not due to inconsistent assignment of categories to responses or to the necessity of memorizing individual stimuli early in training. Instead, we interpret these results as evidence that error-driven associative learning is involved in visual object discrimination and categorization by people, just as in the case of nonhuman animals (Soto & Wasserman, 2010).

If our view is correct, then these results add to a growing body of evidence suggesting that the mechanisms of object recognition and categorization are similar across different species. Prior research has found that both humans and pigeons rely on nonaccidental properties for three-dimensional shape identification (Gibson, Lazareva, Gosselin, Schyns, & Wasserman, 2007). Here, we found that another important aspect of visual categorization is also similar in humans and pigeons: Across species, prediction error seems to be necessary for categorization learning to occur.

The idea that associative-learning principles underlie categorization is not new (Gluck & Bower, 1988; Mackintosh, 1995; Shanks, 1991). However, ours is the first evidence of the involvement of these principles in the categorization of real objects in photographs. These results were obtained using stimuli of greater external validity than the stimuli usually employed in artificial-categorization research. Our study thus joins previous work documenting the importance of associative-learning processes in cognition (for a review, see Siegel & Allan, 1996) and underscoring the generality of learning principles across species, stimuli, and experimental paradigms.

Although we have tried to eliminate alternative explanations of our data, other possible interpretations might be evaluated in the future. For example, it is possible that our pseudocategorization task does indeed foster a memorization strategy that control participants reassessed in Experiment 3 after exposure to new stimuli. Nevertheless, it is unclear why participants in the blocking groups did not reassess their strategy in the testing phase, during which they were presented with novel stimuli and explicitly trained on the categorization task. This and other hypotheses surely merit future empirical testing, but an explanation in terms of error-driven learning remains the most complete and parsimonious account of the available evidence.

In sum, our study provides evidence implicating associative-learning processes in humans' discrimination and categorization of objects in natural scenes. Similar evidence in animals suggests that common learning processes are deployed across diverse species to solve visual categorization tasks. Our study also proves that it is possible to investigate the learning mechanisms involved in object categorization using complex natural stimuli—this is a kind of task that has not been pursued in the past but may represent an interesting line of future investigation.

Acknowledgments

Funding

This research was supported by National Institute of Mental Health Grant MH47313 and by National Eye Institute Grant EY019781.

References

- Ashby, FG. Multidimensional models of categorization. In: Ashby, FG., editor. Multidimensional models of perception and cognition. Hillsdale, NJ: Erlbaum; 1992. p. 449-483.
- Bhatt RS, Wasserman EA, Reynolds WF, Knauss KS. Conceptual behavior in pigeons: Categorization of both familiar and novel examples from 4 classes of natural and artificial stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*. 1988; 14:219-234.
- Cheng PW. From covariation to causation: A causal power theory. *Psychological Review*. 1997; 104:367-405.
- Frank, MJ.; Daw, ND., editors. *Cognition*. Vol. 113. 2009. Reinforcement learning and higher level cognition.
- Gibson BM, Lazareva OF, Gosselin F, Schyns PG, Wasserman EA. Nonaccidental properties underlie shape recognition in mammalian and nonmammalian vision. *Current Biology*. 2007; 17:336-340. [PubMed: 17275301]
- Gluck MA, Bower GH. From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*. 1988; 117:227-247. [PubMed: 2971760]
- Herrnstein RJ. Levels of stimulus control: A functional approach. *Cognition*. 1990; 37:133-166. [PubMed: 2269005]
- Hummel JE. Complementary solutions to the binding problem in vision: Implications for shape perception and object recognition. *Visual cognition*. 2001; 8:489-517.
- Hummel JE, Stankiewicz BJ. Two roles for attention in shape perception: A structural description model of visual scrutiny. *Mechanisms of Visual Attention: A Cognitive Neuroscience Perspective*. 1998; 5:49-79.
- Kamin, LJ. Selective association and conditioning. In: Mackintosh, NJ.; Honig, WK., editors. *Fundamental issues in associative learning*. Halifax, Nova Scotia, Canada: Dalhousie University Press; 1969. p. 42-64.
- Lazareva, OF.; Wasserman, EA. Categories and concepts in animals. In: Menzel, R., editor. *Learning and memory: A comprehensive reference: Vol. 1. Learning theory and behavior*. Oxford, England: Elsevier; 2008. p. 197-226.
- Mackintosh NJ. Categorization by people and pigeons: The 22nd Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*. 1995; 48A:193-214.
- Maia TV. Reinforcement learning, conditioning, and the brain: Successes and challenges. *Cognitive, Affective, & Behavioral Neuroscience*. 2009; 9:343-364.
- Niv Y. Reinforcement learning in the brain. *Journal of Mathematical Psychology*. 2009; 53:139-154.
- Nosofsky RM. Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1984; 10:104-114.
- Nosofsky RM, Kruschke JK, McKinley SC. Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1992; 18:211-233.
- Palmeri TJ, Gauthier I. Visual object understanding. *Nature Reviews Neuroscience*. 2004; 5:291-303.
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical conditioning II: Current theory and research*. New York, NY: Appleton-Century-Crofts; 1972. p. 64-99.
- Riesenhuber M, Poggio T. Hierarchical models of object recognition in cortex. *Nature Neuroscience*. 1999; 2:1019-1025.
- Rumelhart, DE.; McClelland, JL., editors. *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press; 1986.
- Schultz W, Dickinson A. Neuronal coding of prediction errors. *Annual Review of Neuroscience*. 2000; 23:473-500.
- Schwartz, B.; Wasserman, EA.; Robbins, SJ. *Psychology of learning and behavior*. New York, NY: Norton; 2002.

- Serre T, Oliva A, Poggio T. A feed forward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences, USA*. 2007; 104:6424–6429.
- Shanks DR. Categorization by a connectionist network. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1991; 17:433–443.
- Siegel S, Allan LG. The widespread influence of the Rescorla-Wagner model. *Psychonomic Bulletin & Review*. 1996; 3:314–321.
- Soto FA, Wasserman EA. Error-driven learning in visual categorization and object recognition: A common elements model. *Psychological Review*. 2010; 117:349–381. [PubMed: 20438230]
- Vogel EH, Castro ME, Saavedra MA. Quantitative models of Pavlovian conditioning. *Brain Research Bulletin*. 2004; 63:173–202. [PubMed: 15145138]

Fig. 1.

Predicted and observed mean proportions of correct test trials as a function of group (blocking or control) and trial type (training stimuli or novel stimuli). The graphs show the predictions of our model involving error-driven learning (a; for a description of the model and the parameters used in this simulation, see Soto & Wasserman, 2010), replotted data from a previous experiment studying pigeons (b; Soto & Wasserman, 2010), and data from the three experiments reported here (c–e). Error bars represent standard errors of the mean.

Fig. 2.

Examples of stimuli used in the present study and in our previous pigeon experiments (Soto & Wasserman, 2010) on natural-image categorization.

Table 1

Conditions Used in Experiments 1 Through 3

Condition	Training		Generalization testing
	Phase 1: pseudocategorization	Phase 2: categorization	
Blocking—Experiments 1, 2, and 3	Category 1—Set A, Response 1 Category 2—Set A, Response 2 Category 1—Set B, Response 2 Category 2—Set B, Response 1	Category 1—Set A, Response 1 Category 2—Set A, Response 2	Phase 2 training trials Category 1—testing set Category 2—testing set
Control—Experiment 1	—	Category 1—Set A, Response 1 Category 2—Set A, Response 2	Phase 2 training trials Category 1—testing set Category 2—testing set
Control—Experiment 2	Category 1—Set A, Response 2 Category 2—Set A, Response 1	Category 1—Set A, Response 1 Category 2—Set A, Response 2	Phase 2 training trials Category 1—testing set Category 2—testing set
Control—Experiment 3	Category 1—Set B, Response 1 Category 2—Set B, Response 2 Category 1—Set C, Response 2 Category 2—Set C, Response 1	Category 1—Set A, Response 1 Category 2—Set A, Response 2	Phase 2 training trials Category 1—testing set Category 2—testing set

Note: A “set” refers to 10 different photographs depicting an object from a particular basic-level category.