Rules and Resemblance: Their Changing Balance in the Category Learning of Humans (*Homo sapiens*) and Monkeys (*Macaca mulatta*)

Justin J. Couchman,
Department of Psychology, University at Buffalo, the State University of New York

Mariana V. C. Coutinho, and
Department of Psychology, University at Buffalo, the State University of New York

J. David Smith
Department of Psychology and Center for Cognitive Science, University at Buffalo, the State University of New York

Abstract

In an early dissociation between intentional and incidental category learning, Kemler Nelson (1984) gave participants a categorization task that could be performed by responding either to a single-dimensional rule or to overall family resemblance. Humans learning intentionally deliberately adopted rule-based strategies; humans learning incidentally adopted family-resemblance strategies. The present authors replicated Kemler Nelson's human experiment and found a similar dissociation. They also extended her paradigm so as to evaluate the balance between rules and family-resemblance in determining the category decisions of rhesus monkeys. Monkeys heavily favored the family-resemblance strategy. Formal models showed that even after many sessions and thousands of trials, they spread attention across all stimulus dimensions rather than focus on a single, criterial dimension that could also produce perfect categorization.

Keywords

concepts; categorization; rules; family resemblance; comparative cognition; primate cognition; monkeys

Forming categories is an essential function of cognition for humans and nonhuman animals (hereafter animals), and for this reason categorization is a sharp focus of research on humans (Ashby & Maddox, 1992; Murphy, 2002; Rosch & Mervis, 1975; Smith & Minda, 1998) and animals (Cerella, 1979; Herrnstein, Loveland, & Cable, 1976; Medin, 1975; Pearce, 1994; Smith, Minda, & Washburn, 2004; Wasserman, Kiedinger, & Bhatt, 1988). One goal of this article is to produce a synergy between these research areas by extending to the study of animal categorization insights from research on human categorization.

In particular, research indicates that humans have multiple systems of categorization that emphasize different organizational principles for categories and that may have different organizational bases and loci in brain. Humans have the capacity to spread attention broadly over several stimulus dimensions, and to use the features that are probabilistically shared by members of a category to recognize its members. Following philosophical work by Wittgenstein (1953/2001) and grounding psychological work by Rosch (e.g., Rosch & Mervis, 1975), this overall-similarity approach to categorization can be called family-resemblance (FR) categorization to capture the idea that members of a category family are generally featurally similar. Related ideas are formalized in the animal literature's linear feature model (e.g., Huber, 2001, Jitsumori, 1994). FR learning is often studied using polymorphous categories in which several stimulus features are partially informative about category membership (Cook & Smith, 2006; Huber & Lenz, 1993; Jitsumori, 1994, 1996; von Fersen & Lea, 1990; Lea & Willis, 2008), and this practice will be followed here. One goal of the present research is to evaluate the importance of FR category learning—relative to other approaches—in nonhuman primates' overall categorization competence.

Humans also have the capacity to focus attention sharply toward a single stimulus dimension, in essence using a unidimensional rule for grouping stimuli into categories. Research demonstrating this has used categories that are defined by a 100%-diagnostic (criterial) attribute (Ahn & Medin, 1992; Ashby & Ell, 2001; Erickson & Kruschke, 1998; Nosofsky et al., 1994; Regehr & Brooks, 1995). Interest in single-dimensional discriminations or featural rules is also traditional in animal research (Mackintosh, 1975; Sutherland & Mackintosh, 1971), as when Krechevsky (1932) asserted that rats test "hypotheses" or when Lashley (1942) found that rats sometimes made cue-response associations discontinuously. In fact, the multidimensional discrimination task used by Krechevsky and others has the same structure as the criterial-attribute (CA) category task used in human research (Smith et al., 1993). Given the lasting interest in this issue, the present research takes a new approach toward analyzing the balance—in primates' overall categorization system—between learning a single-dimensional criterial attribute and appreciating the overall family resemblance among category members.

FR and CA category learning have been clearly dissociated in humans (Kemler Nelson, 1984; Love, 2002; Smith & Shapiro, 1989; Smith et al., 1993; Waldron & Ashby, 2001). Generally speaking, research links FR category learning to incidental learning conditions and procedural learning processes. Research links CA category learning to intentional learning conditions and explicit, usually verbally mediated cognition. For example, Smith et al. (1993) found that adults suffering from depression performed more poorly than nondepressed individuals on CA but not FR tasks. Smith et al. suggested that humans with depression still have the ability to categorize using a passive and procedural fallback strategy even when they do not have available the cognitive resources necessary for intentional, rule-based categorization. Kemler Nelson (1984) and Smith and Shapiro (1989) used category tasks that allowed FR or CA learning to show that intentional and incidental learning conditions, respectively, favored CA and FR learning. Waldron and Ashby (2001) found that concurrent Stroop interference selectively impaired performance in CA tasks compared to FR tasks. Love (2002) found that incidental learning conditions selectively impaired CA tasks compared to FR tasks. It is also known that young, mentally retarded, and impulsive children tend to favor FR category-learning strategies, probably due to the less strategic character of their immature cognition (Kemler, 1982; Kemler Nelson & Smith, 1989; Shepp, Burns, & McDonough, 1980; Smith & Kemler Nelson, 1988; L.B. Smith & Kemler, 1977; Ward, 1983). It is an interesting possibility—with cross-species implications—that FR categorization is a more primitive, fallback strategy that is possible in the absence of more sophisticated cognitive strategies. And it is an interesting question whether nonhuman primates have as robust a capacity for deliberate and analytic cognition as humans. In fact, if CA category learning is

closely linked in humans to declarative consciousness or to verbal mediation, then monkeys might prove to be limited in using single-dimensional rules or criterial attributes in categorization, and the balance might be shifted toward FR strategies in their category learning.

Recently, Ashby and colleagues (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby, Queller, & Berretty, 1999; Ashby & Ell, 2001; Ashby & O'Brien, 2005) grounded the CA-FR distinction in neurobiological evidence. They hypothesize that performance in CA tasks is generally mediated by frontal cortical structures (e.g., prefrontal cortex and the anterior cingulate gyrus) that may also serve executive attention and that would be suited to managing hypothesis testing and switching. They hypothesize that FR tasks are localized in areas of the striatum and basal ganglion and are less dependent on executive attention, hypothesis testing, and verbal mediation.

That humans have these multiple systems for organizing categories may be a general testimony to the adaptive significance of forming categories and behavioral equivalence classes, and to the need for having multiple, complementary solutions for these kinds of cognitive problems. However, this contrast between the striatal system and the frontal system raises interesting questions about species differences in categorization. Monkeys are likely weaker than humans in the cognitive capacities supported by frontal brain systems. They have proportionally smaller frontal cortices (Semendeferi et al., 2002). They are compromised relative to humans on frontal tasks that involve response competition or response inhibition (e.g., Stroop or Flanker tasks– Roberts, 1996; Washburn, 1994). They might also have a rudimentary frontal categorization system. From this perspective, too, one might predict that the balance would shift toward FR strategies in monkeys' category learning. If so, this would provide an important insight into the character of nonhuman primates' categorization system, and an insight into the ancestral categorization system from which the human categorization system emerged by extension and amplification.

To comparatively explore the CA and FR modes of categorization here, we adopted Kemler Nelson's (1984) innovative and sensitive technique for evaluating the preferential use of CA or FR categorization strategies. To create her dissociation between explicit and implicit strategies, Kemler Nelson used the category structure summarized abstractly in Figure 1. The stimuli were composed of four binary features (stimulus values 0 and 1 represent the two possible values for each dimension). Dimension 1 in the category structure was perfectly diagnostic and criterial. Its sole use as a category rule (0: Category A; 1: Category B) would produce perfect categorization performance. In addition, though, Dimensions 2–4 embodied a parallel differentiating principle for the two categories. Category A members are generally similar across these dimensions (sharing the same dimensional values), as are Category B members. Accordingly, participants could also solve Figure 1's category problem by spreading attention over 3 or even all 4 stimulus dimensions, weighting them all about equally, and using the overall similarity among Category A and Category B members to guide categorization decisions.

Kemler Nelson (1984) made particular use of the test stimuli shown in Figure 1 to ask whether participants were using the single-dimensional rule or the overall-similarity strategy. These stimuli place the dimensional rule and the overall similarity relations into direct conflict. The first test stimulus would be placed in Category A or B, respectively, based on the use of the rule or overall similarity. The second test stimulus would be placed in Category B or A, respectively, based on the rule or overall similarity.

Kemler Nelson (1984) found that human participants strongly relied on the rule-based strategy under intentional learning conditions. However, she found that this strategy was often supplanted by an overall-similarity strategy under incidental or unintentional learning

conditions. The present experiments replicate Kemler Nelson's (1984) original finding with humans, and then use the same category structure to evaluate in the category learning of monkeys the balance between CA or FR strategies.

Experiment 1: Humans

Experiment 1 compared the modal categorization strategies of intentional and incidental human category learners. This gave us a grounding human perspective by replicating Kemler Nelson's (1984) original dissociation using our own stimuli, tasks, and procedures.

Method

Participants—To obtain 36 learners in each condition, 84 participants were recruited. Participants were undergraduates from the University at Buffalo, the State University of New York, who participated in a session lasting about an hour to fulfill a course requirement. Participants were in their late teens or early twenties with apparently normal or corrected-to-normal visual acuity. The approximate racial mix of our participant pool was 61% Caucasian, 25% Asian, 8% African-American, and 6% Other. Fifty six percent of participants were male; 44% were female.

Categorization stimuli—The stimuli were multi-colored circles divided into four 90-degree sectors (Figure 2). Each sector was a dimension in the category structure of Figure 1 and could take on one of two possible colors. Though Kemler Nelson (1984) used facial features rather than colors, the logical category structure was identical to hers.

In the example shown in Figures 1,2, each 0–1 value in Figure 1 represents one of two possible colors for a disk sector in Figure 2. For example, a 0 or 1 in the leftmost position in Figure 1, respectively, might represent a brown or pink 3-o'clock disk sector in Figure 2. In the figures, Figure 1, the leftmost stimulus dimension and the 3-o'clock disk sector is defining, because 0 (brown) occurs only and always in Category A, whereas 1 (pink) occurs only and always in Category B. In addition, every Category A exemplar is composed predominantly of 0 values (green, white, blue), whereas every Category B exemplar is composed predominantly of 1 values (magenta, yellow, gray). The first (top) stimulus in each category in each figure is the prototype for that category. It contains all four features that are characteristic of its category, and thus it perfectly exemplifies both the criterial-attribute rule and the family-resemblance structure for that category.

To ensure that no particular color or dimensional position influenced the outcome, we created 3 alternative color sets and 4 alternative dimensional structures for the categories. Color Sets 1 and 2 had exactly opposite colors – that is, a 0 in Color Set 1 would be a 1 in Color Set 2. Color Set 3 was an alternative randomization of the colors used in Color Sets 1 and 2. The 4 dimensional structures were created by shifting the values in Figure 1 (which shows Dimensional Structure 1) one position to the right and then moving the last value to the beginning of the string. So, for example, Category A in Dimensional Structure 2 would have the structure "0000, 0010, 0001, 1000", making the second position of each exemplar the defining feature. Test stimuli were shifted in the same way, so that each dimensional structure had the same logical structure.

In all stimulus sets, the test exemplars, like those illustrated in Figure 2, contained the criterial attribute of one category opposed by the family-resemblance structure of the other category along the noncriterial dimensions. Thus, a participant using a CA categorization strategy would put the top and bottom test stimuli into Category A and B, respectively. A participant using an FR categorization strategy would put the top and bottom test stimuli in Category B and A, respectively.

Categorization trials—Each trial consisted of a to-be-categorized disk presented in the top center of a 17.5-inch computer screen against a black background. Below each disk appeared screen icons ('A' or 'B') reminding participants about their response options. These responses were selected by choosing labeled keyboard keys that reflected the spatial layout of the response icons on the screen. Participants received a point and heard a 0.5 s computer-generated reward whoop for each correct response. They heard an 8-s computer-generated penalty sound and lost a point for each error. In addition, after each trial, participants received a scorecard textbox on the screen that gave them a +1 or −1 for that trial and told them how many points they currently had. After the feedback, the screen cleared and the next trial was presented.

Procedure: intentional learning condition—Participants were placed randomly into intentional and incidental conditions. In the intentional condition, participants were told that they would see many multi-colored circles and should figure out which were members of Category A or B. They were then given 80 training trials, divided into 10 8-trial blocks, with each block containing all the category members from both categories (but no test stimuli). After these training trials, no feedback was given for the rest of the experiment. Participants were told to continue categorizing as they had learned to do. The participants then received 3 stages of transfer as follows. First, there were 3 blocks of 10 trials, with each block containing one each of the Category A and B stimuli and one extra instance of each prototype (the first stimulus in each category). Second, there were 2 blocks of 18 stimuli, with each block containing each of the category stimuli, 2 extra prototypes, and 4 instances of each of the test stimuli. Third, there were 3 blocks of 18 stimuli, with each block containing all 16 stimuli possible in the four-dimensional stimulus space and an extra prototype from each category. This third set of blocks was used to determine whether the participants were tracking a single, but noncriterial, attribute.

Procedure: incidental learning condition—The incidental condition was identical except that in training participants were given a memory cover task. That is, on each trial they were shown 2 stimuli, one from each category, and simply asked which they had seen more often. They were warned that as they progressed they would see many repeats. Each stimulus was placed over an A or a B, which the participant perceived as being merely for response purposes (i.e. "push 'A' if you have seen the circle over the 'A' more often"). Participants were given 10 blocks of 4 trials each, with each block containing all 4 Category A stimuli and all 4 Category B stimuli in random pairs. They received a whoop and a point for each answer. After training, they were told that the circles over the 'A' were members of Category A, and that the circles over the 'B' were members of Category B, and that they were to continue categorizing the stimuli in that way. They, too, received no feedback in the 3 transfer phases that were exactly as already described.

Results

Levels of learning—In the intentional and incidental conditions, respectively, 36 of 40 (90%) and 36 of 44 (81%) participants were defined as category learners by achieving at least 70% correct responses to the eight category stimuli. These percentages of learners suggest an advantage for the intentional condition. From this point forward, the data from 4 and 8 nonlearners in the intentional and incidental conditions, respectively, were excluded from analysis so as to focus on the data from the 36 learners in each condition. The 36 learners in each condition represented 3 counterbalances of the 3 color sets and the 4 dimensional structures already described.

The intentional condition's advantage was confirmed by comparing the errors made during training. Learners in the intentional condition averaged 4.3 errors ($SD = 7.12$) on the 80 training

stimuli, learners in the incidental condition averaged 9.4 errors ($SD = 7.05$), $t(70) = 3.06$, $p < .05$. Intentional learners mastered the categories more easily.

Types of learning: test stimuli—We first examined the modal categorization strategy for the two conditions using the two test stimuli that were presented 22 times (11 times each) to each participant. The CA responses to the test stimuli averaged 14.6 ($SD = 9.38$) and 10.5 ($SD = 7.85$) in the intentional condition and incidental conditions, respectively, $t(70) = 2.0$, $p < .05$. Just as in Kemler Nelson's (1984) original experiment, intentional and incidental learners produced more and fewer CA responses on the test stimuli.

In the intentional condition, 21 of 36 participants produced a strong CA pattern (18 to 22 CA responses of the maximum 22 possible), averaging 21.9 CA and 0.1 FR responses to test stimuli. Six learners produced a mixed CA-FR response pattern, averaging 10.3 and 11.7 CA and FR responses. Nine learners produced a strong FR pattern (0 to 4 CA responses out of 22), averaging 0.4 and 21.6 CA and FR responses. Humans clearly produce a mixture of category-learning strategies even under intentional learning conditions. These differences are potentially important theoretically. Evidently humans are not locked into a monolithic CA approach to tasks of this kind. This raises a stronger prospect of cross-species continuity with monkeys than if humans produced a rigid, invariant rule-based performance.

In the incidental condition, humans once again produced a mix of strategies, but now FR responding was more prominent. Thirteen of 36 participants produced a strong FR response pattern, averaging 1.8 and 20.2 CA and FR responses on the 22 test stimuli. Two participants (averaging 15.5 FR responses on the test stimuli) were less strong FR learners. Seven participants (averaging 11.3 FR responses) produced a mixed CA-FR pattern. Five participants (averaging 16.2 CA responses) produced a weak CA data pattern. Nine participants (averaging 20.8 CA responses) produced a strong CA data pattern.

Following previous researchers (Kemler Nelson 1984, 1988; Ward & Scott 1987; Smith & Shapiro 1989), we performed additional analyses, using the blocks in which all 16 possible stimuli were presented, to evaluate whether some participants were obeying systematically the information provided by some single, noncriterial attribute. This obedience would let them perform sufficiently well to be considered a learner, but it would also create a deceptive and misleading FR response pattern on the test stimuli. We found (as also found in previous research) that no participant in the intentional condition, and only one participant in the incidental condition, tracked a noncriterial attribute. Thus, the incidental FR learners can be correctly considered as participants who used a strategy fundamentally different from the intentional CA learners. They were not merely participants who used a non-ideal single-dimensional strategy.

These results represent a clear replication of Kemler Nelson's (1984) results. This is theoretically important because both the stimuli and the cover task (in the incidental condition) were different than Kemler Nelson (1984), and are arguably more abstract. This demonstrates that the dissociation she found is robust to these differences. This replicable dissociation in humans grounds the attempt to evaluate the balance in monkeys between CA and FR category-learning strategies.

The preceding analyses were somewhat incomplete because they only analyzed humans' responses to the two critical-test stimuli. They do not let one assess humans' overall attentional strategy in the task. Many participants' responses to the test stimuli indicated mixed or weak strategies in the direction of FR or CA learning. It is not clear, without an attentional profile, what type of strategy caused them to categorize the test stimuli in this way. Accordingly, to

determine that overall strategy, we adopted an alternative analytic technique that involved formal modeling to estimate the overall attentional profile of each participant.

A Formal-Modeling Perspective

We sought converging evidence for Experiment 1's conclusions by also analyzing the eight stimuli in the two training categories (Figure 1). For this purpose, data from both category-learning conditions were also modeled mathematically using a multiplicative prototype model with an additional guessing parameter. This model was also used in Smith and Minda (1998, 2000) and is described fully there. The choice of this model was natural given category structures in the present experiments that were linearly separable, with a good family-resemblance structure, without exceptions that might require dedicated exemplar-memorization processes, and that could be well learned by humans through a process like prototype abstraction. The choice of this model does not presume that prototype abstraction is the norm for human categorization, or that other representational formats for categories are to be discounted. The present perspective on modeling was undertaken only to gauge humans' attentional processes in these categorization tasks, and the prototype model was adequate to this task and unbiased regarding this task.

To instantiate this model, we supposed that each to-be-categorized item would be compared with the category prototype (e.g., 0000 for category A or 1111 for category B in Figure 1) along the four independent stimulus dimensions. This comparison yielded a measure of psychological distance between the prototype (the category representation) and the to-be-categorized stimulus. Matching features contributed 0.0 distance, while mismatching features contributed the amount of their dimension's attentional weight. Our measure of psychological distance incorporated attention that was presumed to be limited and distributed across the stimulus dimensions. That is, the attentional weights given to the four dimensions were constrained to sum to 1.0. So, for example, if a participant allocated attention homogenously (.25 to each dimension), and a stimulus shared 3 features in common with the prototype, the psychological distance would be 0.25. Or, if a participant allocated all of attention to one stimulus dimension, as would be the case if they were following a rule, then all stimuli fulfilling the criteria for that rule would be judged to have a distance from the prototype of 0.0.

The modeling translated psychological distance into a measure of psychological similarity by mathematically inverting distance (more distance, less similarity) by making similarity an exponential-decay function of distance. A sensitivity parameter governed this inversion:

$$similarity = e^{-sensitivity \times distance}.$$

Then, our measure of endorsement level came from entering similarity into a choice rule with this general form:

$$endorsement\ level_{cat1} = similarity_{cat1} / (similarity_{cat1} + similarity_{cat2}).$$

We also assumed that some proportion of the time participants, especially in the early stages of categorization, simply guessed category A or B haphazardly, and so a guessing parameter was added to let the model incorporate this aspect of categorization behavior.

We used standard hill-climbing methods to find the parameter settings that let the model reproduce as well as it could a categorization performance profile (i.e., the endorsement levels given to 8 training stimuli). We used the sum of the squared deviations between the observed and predicted endorsement levels to find the best fit of the model to the observations. The

model hill climbed toward the best-fitting configuration by minimizing the sum of the squared deviations between the predicted probabilities and the participant's observed performance. The analysis of CA and FR responses described above only took the two test stimuli into account. The attentional analyses offered by the model—based on four times as many stimuli—are therefore more systematic and comprehensive.

Humans: intentional condition—Figure 3A shows the attentional weights estimated for each stimulus condition for each participant in the intentional condition. Each column of weights along the X-axis represents one participant's allocation of attention across the four dimensions in the task. Data points at 100% reflect perfectly focused attention to the task's criterial attribute and are a strong indication that the participant was using the CA learning strategy. The average attentional weight given to the CA in the intentional condition was 72.6%. Twenty-three of 36 learners allocated—as estimated by the model—100% attention to the CA. These results differ slightly from the test-stimulus analysis above, because now the eight training stimuli were the focus of analysis. Yet the two analytic approaches were clearly mutually confirming. Intentional participants made 66.3% CA responses to test stimuli and gave 72.6% attention to the CA in classifying training stimuli. The two analyses found that 21 of 36 participants, or 23 of 36 participants, followed a strong CA strategy. The attentional estimates from the training-stimulus modeling would be expected to be less variable and more powerful given that they are based on more stimuli and a larger sample of each participant's performance.

Humans: incidental condition—Figure 3B shows the attentional weights estimated for each stimulus dimension for each participant in the incidental condition. Each column of weights along the X-axis represents one participant's allocation of attention across the 4 dimensions in the task. Data points at 100% reflect perfectly focused attention to the task's criterial attribute or to a noncriterial attribute and are a strong indication that the participant was using a rule-based learning strategy. The average attentional weight given to the CA in the incidental condition was 39.6%, significantly less than in the intentional condition, $t(70) = 4.2$, $p < .05$. Only 2 of 36 learners allocated 100% attention to the CA. The model also correctly discovered the one case of a participant tracking a noncriterial attribute.

Once again the test- and training-stimuli analyses were mutually reinforcing. Incidental participants made 47.7% CA responses to test stimuli, and gave 39.6% attention to the CA as estimated by the model. The two analyses found that 9 of 36 participants, or 2 of 36 participants, followed a strong CA strategy.

Experiment 2: Monkeys

Method

Participants—Murph (13 years old) and Lou (13 years old) were tested. They had been trained, using procedures described elsewhere (Rumbaugh, Richardson, Washburn, Savage-Rumbaugh, & Hopkins, 1989; Washburn & Rumbaugh, 1992), to respond to computer-graphic stimuli by manipulating a joystick. They had also been tested in prior studies on a variety of computer tasks, including dot-distortion categorization, numerical discriminations, uncertainty monitoring, same-different judgments, and visual search. The monkeys had never before encountered these color-disk stimuli or this task structure.

These two animals have been highly accurate experimental participants in many procedures. We believe that they should be especially strong candidates for revealing CA responding in a task of this kind. We also believe that their long experience could have biased them somewhat toward finding CA task solutions. Thus, we believe that the present subjects and experiments represent, if anything, conservative tests of whether monkeys generally will favor CA or FR

responding. We acknowledge, however, that there is some chance that we selected two inherently FR-preferring monkeys.

The monkeys were tested in their home cages at the Language Research Center of Georgia State University, with *ad lib* access to the test apparatus, working or resting as they chose during long sessions. Each monkey continued working in the task until a sufficient number of trials had been collected (approximately 11,000). The animals were neither food deprived nor weight reduced for the purposes of testing and they had continuous access to water.

Apparatus—The monkeys were tested using the Language Research Center's Computerized Test System—LRC-CTS (described in Richardson et al., 1990; Rumbaugh et al., 1989; Washburn & Rumbaugh, 1992)—comprising a Compaq DeskPro computer, a digital joystick, a color monitor, and a pellet dispenser. Monkeys manipulated the joystick through the mesh of their home cages, producing isomorphic movements of a computer-graphic cursor on the screen. Contacting appropriate computer-generated stimuli with the cursor brought them a 94-mg fruit-flavored chow pellet (Bio-Serve, Frenchtown, NJ) using a Gerbrands 5120 dispenser interfaced to the computer through a relay box and output board (PIO-12 and ERA-01; Keithley Instruments, Cleveland, OH). Correct responses were accompanied by a computer-generated whooping sound that bridged the animals to their reward. On incorrect responses, the screen froze with the wrong response visible, and there was a computer-generated buzzing sound and a timeout that lasted 20 s.

Training—Monkeys were trained using blocks of 10 trials, with each block containing the eight training stimuli and an additional instance of each prototype until they achieved 80% overall accuracy. During training, they received a whoop and a food pellet for every correct answer, and a buzz and a 20-s timeout for every incorrect answer.

Testing—It was important not to reinforce the test stimuli differentially depending on the animals' response to them, and we thought it would be ideal not to reinforce the test stimuli directly at all. This feedback could bias animals toward one response strategy or the other. However, we could not withhold feedback completely during testing as we did with humans. Our compromise was to adapt the animals to receiving deferred feedback in the task after blocks of trials. They were trained to respond to 6 stimuli before receiving a group of 6 outcomes, with all of the correct-response whoops and food pellets delivered first, followed by all the incorrect-response buzzes and timeouts. Under this reinforcement regimen, it was possible to reinforce responses to the test stimuli indirectly and non-differentially. We rewarded positively either response to all test stimuli, and these outcomes when they occurred were simply added as a reward during each deferred group of 6 outcomes.

We trained the animals to tolerate this deferred, rearranged feedback using 21 gradually increasing levels of deferral. We used cumulative block-length probabilities that went from 100 0 0 0 0 0 (all one-trial blocks), through 40 65 85 100 100 100 (40% one-trial blocks, 25% two-trial blocks, etc.), through 0 5 25 45 70 100 (0% one-trial blocks, 5% two-trial blocks, etc), and finally to 0 0 0 0 0 100 (all six-trial blocks). We moved to the next level in this table each time animals' performance recovered to 80% at the new level. In this way we gradually weaned animals toward less immediate and more deferred reinforcement. By enforcing 80% accuracy at each level, we ensured that the macaques treated each trial as valid even though they were not directly reinforced for it.

We believe that this method of deferred feedback could contribute to other lines of comparative research. The contingencies regarding reward and punishment for responses made in different stimulus contexts are typically made clear to animals through direct and immediate feedback signals. However, this methodology always supports and encourages a particular associative

theoretical framework for explaining animal behavior. It may also condition animals' response patterns in a fine-grained manner that prevents one from seeing the animals' own decisional and choice processes at work. By removing immediate reinforcement from the experimental context, one may let animals reveal their own independent choice and decision processes more clearly. In our case, deferred feedback allowed us to determine their categorization choices on test stimuli without differentially reinforcing any particular categorization strategy.

The adaptation to deferred reinforcement occurred after animals reached 80% accuracy in training, but before they moved on to the testing phase of the experiment. Monkeys adapted to deferred feedback easily, with each achieving 80% accuracy in fewer than 4 sessions.

During the testing phase, monkeys received blocks of 12 stimuli, with each block containing the 8 category stimuli, 1 extra instance of each prototype, and 1 instance of each test stimulus. All trials during the testing phase were presented in 6-trial blocks that received deferred and rearranged reinforcement. Each session, after the macaques had graduated to 6-trial deferred reinforcement, started with a brief training period as described above, with no test instances presented, and then transitioned immediately to 6-trial blocks with deferred reinforcement and with test instances presented.

Results

Levels of learning—Monkeys generally showed the highest levels of performance during training trials when they received trial-by-trial reinforcement. Murph and Lou each reached over 80% correct during training trials, as this was the performance needed to progress through the experiment which both animals did. Performance declined somewhat during testing trials, as the animals completed 6-trial blocks with deferred feedback. Murph and Lou, respectively, averaged 70% and 72% correct during testing trials.

Type of learning: Murph—Murph exhibited a weak FR response pattern. Of the 948 test stimuli he was shown, he made 610 FR responses, compared to only 338 CA responses. Murph clearly did not discover the single-dimensional rule that was embodied in the task's criterial attribute.

To confirm this pattern, we modeled using the method already described Murph's performance on the eight training stimuli. The model estimated that, over all sessions, Murph averaged 28% attentional weight to the criterial attribute compared to 20%, 15%, and 37% to the noncriterial attributes. The model estimated that Murph never averaged over 45% attentional weight to the CA in any session, and in fact attention to the CA tended to decrease over sessions. This broad and homogeneous distribution of attention clearly shows that that he was following an FR categorization strategy. Figure 4A shows the estimates of Murph's four attentional weights for the four stimulus dimensions over trial blocks during a representative session. The figure confirms that Murph was using a broad attentional strategy that was distributed about equally across the four stimulus features. (During this session, Murph averaged 27% attentional weight to the criterial attribute compared to 23%, 12%, and 38% to the noncriterial attributes.)

Type of learning: Lou—Of 1138 test stimuli, Lou made 368 FR responses and 770 CA responses. Lou also did not fully discover the category task's criterial attribute, though he did make more CA responses than FR responses and more CA responses than Murph did.

However, the formal modeling that incorporates performance on all eight stimuli, and not just on these two test stimuli, suggests that Lou's CA responding was actually weaker than it appeared to be for the test stimuli. Lou's performance on the eight training stimuli was modeled as already described. The model estimated that, across all sessions, Lou averaged 36% attentional weight to the criterial attribute, compared to 12%, 32%, and 20% to the noncriterial

attributes. The model estimated that Lou never averaged over 40% attentional weight to the CA in any session, and in fact attention to the CA tended to decrease over sessions. This broad and homogeneous distribution of attention clearly suggests that he was following an FR categorization strategy, though with more attention given to the CA. Figure 4B shows the estimates of Lou's four attentional weights for the four stimulus dimensions over trial blocks during a representative session. The figure confirms that Lou was using a broad attentional strategy that was distributed about equally across the four stimulus features. (During this session, Lou averaged 35% attentional weight to the criterial attribute, compared to 6%, 39%, and 20% to the noncriterial attributes.)

Experiment 3: Monkeys

Experiment 2 gave the monkeys an intensive, extended exposure to one category task and structure. Given this exposure, neither animal learned to tune sharply to the task's single-dimensional or rule-based solution. Harlow (1949) took an alternative approach to the problem of eliciting rule use from animal subjects. He gave monkeys long series of simple rule-use problems, so that they would have many successive opportunities to engage in rule discovery. Given this exposure, he did find that they had some capacity to learn rules, apply them, and to change their solution quickly. Therefore, we adopted Harlow's approach in Experiment 3. We gave the monkeys a randomly selected different set of stimulus colors and a randomly selected different dimensional structure (i.e., with varying criterial attributes each day), so that new category learning had to occur each day, including either the appreciation of a new criterial attribute or a new family-resemblance organization. And so we asked whether animals might acquire a CA learning set under these conditions, even though they had not under the extended training in Experiment 2.

We also modified the reinforcement regimen from that in Experiment 2. In Experiment 2, the deferred and rearranged feedback may have made ongoing category learning and rule discovery difficult. Therefore, in Experiment 3 we adopted a reinforcement regiment that rewarded many trials directly and that should have made ongoing learning and rule discovery during testing feasible for the animals.

Method

Participants, apparatus—Murph and Lou were tested using the same paradigm as Experiment 2.

Training—Training was like that in Experiment 2, except that now a different reinforcement regimen was applied. We rewarded 75% of trials directly, using the whoops, food pellets, buzzes, and timeouts already described. We provided no feedback on 25% of trials, for reasons to be described momentarily. On these trials, moving the cursor to A or B immediately produced the next trial. As in Experiment 2, they trained until they reached the learning criterion of 80% correct.

Testing—Testing differed from Experiment 2 only in the reinforcement regimen. During testing, the critical test stimuli (2 in every 12 trials—16.6% of trials) were never rewarded. In addition, 10% of category stimuli were also not rewarded. In combination, these 26.6% of non-rewarded trials nearly matched the 25% non-rewarded trials to which the animals were already accustomed. This made for a seamless transition from training to testing, made feedback much more direct (supporting ongoing category learning and possibly rule discovery), but still allowed the test stimuli to never be directly or differentially rewarded.

Sessions—At the beginning of each session, one of four dimensional structures, each with a different defining category rule, was randomly selected. Eight different colors (two possible

colors for each disk sector) were randomly selected. Thus, each day the CA dimension and the CA color were highly changeable, as was the nature of the family resemblance within each category. Each monkey completed 10 sessions (10 category tasks). In each case, after completing training, each monkey completed exactly 120 testing blocks (1440 trials).

Results

Levels of learning—Muph and Lou averaged 92.3% and 85.2% correct, respectively, during testing trials. This increased level of learning compared to Experiment 2 indicates that the task and reinforcement structure was indeed advantageous for the monkeys.

Type of learning: Murph—Once again, Murph exhibited a weak FR response pattern. Of the 2400 test stimuli he was shown, he made 1678 FR responses, compared to only 722 CA responses. Although Murph learned quickly and reached a high level of accuracy, he clearly did not discover the single-dimensional rules available within each task.

Murph's performance data on the training stimuli were modeled as already described, using all of the responses made to the 8 category stimuli during the testing phase (including the unrewarded 10%). The model estimated that Murph averaged 39.8% attentional weight to the criterial attribute overall, compared to 18.4%, 26.7%, 18.5%, and 18.2% to attributes 1–4, respectively, when they were noncriterial. This broad distribution of attention confirms that he was following an FR categorization strategy. Murph followed this strategy throughout his 10 sessions. Figure 5A,B show the estimated attentional weights for Murph's first and last sessions, respectively.

Murph did generally learn to give the CA in a task slightly more weight as he progressed through his sessions. Ignoring Block 1 (during which there was a lot of chance responding absent any robust task knowledge), Murph was estimated to give the CA .29 weight in Block 2, but .48 weight in Block 12. This gradually increasing emphasis on the CA is natural, and still consistent with FR responding, because even on a purely associative basis that attribute is most informative and most highly correlated with the reward structure of the task. But Murph clearly did not attend solely to the CA even though it could support perfect categorization. There was no hint that he was gradually developed a learning set for CA rule discovery.

Type of learning: Lou—Lou also exhibited a weak FR response pattern. Of 2400 test stimuli, he made 1444 FR responses and 956 CA responses. Lou also did not reliably discover the category rules. To the contrary, the balance of his categorization responses shifted away from CA and toward FR in Experiment 3.

Lou's data were also modeled as described earlier. The model estimated that Lou averaged 23.6% attentional weight to the criterial attribute overall, compared to 25.1%, 30.5%, 19.5%, and 26.2% to attributes 1–4, respectively, when they were noncriterial. This broad distribution of attention confirms that he was following an FR categorization strategy. Lou followed this strategy throughout his 10 sessions. Figure 5C,D show the estimated attentional weights for Lou's first and last sessions, respectively.

Unlike Murph, Lou did not generally learn to give the CA more weight as he progressed through his sessions. Ignoring Block 1 again, Lou was estimated to give the CA .33 weight in Block 2, but .36 weight in Block 12. There was clearly no hint that Lou was gradually developed a learning set for CA rule discovery.

General Discussion

Humans replicated Kemler Nelson's (1984) result using abstract stimuli and a different memory cover task. Intentional category learners made significantly more CA responses to the transfer stimuli than did incidental category learners. Formal models—that estimated the attentional weights of both types of learners—confirmed that intentional learners focused attention more sharply on the criterial attribute. Twenty-three of 36 of these participants allocated all of their attention to the criterial attribute compared with only 2 incidental learners. The formal models also extended Kemler Nelson's analytic framework to include the training and test stimuli in the task, not just the test stimuli, providing a more comprehensive analytic framework that could be used to analyze comparatively the monkeys' performance.

In contrast, two rhesus macaques showed little preference for the CA solution in these category tasks, even after thousands of trials that all could have been answered correctly based on one stimulus dimension. Formal models estimated that monkeys' attention was broad and substantially multi-dimensional. Overall, the two monkeys made 1,108 CA responses and 978 FR responses over 2,086 test stimuli in Experiment 2 and 1,678 CA and 3,122 FR responses over 4,800 test stimuli in Experiment 3. Monkeys clearly never adopted a criterial rule. For some reason, monkeys did not recruit the processes and representations that would support perfect performance through rule use. Many intentionally learning humans did.

These results continue to show the value of the general theoretical position that FR categorization is the default mode of approach to tasks when full executive functioning is underdeveloped, inhibited, or lacking for any reason. Once again here, humans often showed FR categorization when they were denied the chance to recruit active, deliberate category-learning strategies. This was shown previously to be the case with depressed individuals (Smith et al., 1993) and in adults operating under a competing working-memory load (Waldron & Ashby, 2001). It is also known that young, mentally retarded, and impulsive children tend to favor FR category-learning strategies based in the use of overall-similarity relations, probably due to the less strategic character of their immature cognition (Kemler, 1982; Kemler Nelson & Smith, 1989; Shepp et al., 1980; Smith & Kemler Nelson, 1988; L.B. Smith & Kemler, 1977; Ward, 1983). It is intriguing theoretically that monkeys here also heavily weighted FR solutions to category problems. It implies that it is an extremely general claim—spanning a substantial portion of the primate lineage at least—that FR categorization emerges as a primitive, fallback strategy when more sophisticated cognitive strategies are lacking for any reason.

The present results also continue to show the value of a multiple-systems perspective on categorization urged by Ashby and colleagues (Ashby et al., 1998, 1999; Ashby & Ell, 2001; Ashby & O'Brien, 2005) and supported in research by Kemler Nelson (1984), Love (2002), Smith & Shapiro (1989), Smith et al., (1993), Waldron & Ashby, (2001), and others. Multiple systems theory holds that there are different representational processes that are used in different situations. The present research supported this idea by showing that two different representational strategies—CA and FR—are favored by intentional and incidental learners, respectively, in a categorization task. Monkeys at least showed a well developed FR categorization strategy. They spread their attention across stimulus dimensions, integrating information across these dimensions, rather than focusing on a single-dimensional, defining rule.

It is interesting to consider why monkeys preferred FR responding in our tasks. One possibility is that CA processing may not be the best strategy in every situation. For example, Jitsumori (1993) found that pigeons were better able to transfer knowledge of learned features to novel stimuli than humans, possibly because they spread their attention across the feature space. It

is also possible that the environment might be changeable, so that the weight that would optimally be given to different features defining a category would change through time. In that case, the organism that had preserved a broad attentional strategy would best preserve its capacity to track these changes and respond to them adaptively. Indeed, the structure of Experiment 3 may have simulated these natural changes. The changing category problems could even have somewhat encouraged FR responding, granting monkeys a consistent approach to these multi-dimensional stimuli that served them well day after day.

Another possibility is that our tasks did not do full justice to monkeys' capacity to use rules in category tasks. Our tasks presented two equally viable solutions in parallel to monkeys. They could adopt either and still achieve a high level of performance. Consequently, there was no direct pressure on them to master the rule aspect of the task. We tapped, therefore, monkeys' processing preference or default tendency in category tasks rather than their ultimate rule-using capacity in category tasks. They might become sharper and more focused rule users under conditions that forced them to do so. In fact, in Smith, Beran, Crossley, Boomer, and Ashby (in press), these same two monkeys did show some capacity to selectively attend to a single dimension given two qualitatively different dimensions in a category-learning task.

Similarly, it has been established that macaques have the ability to attend to the abstract dimension of sameness-difference in auditory (Wright, Shyan, & Jitsumori, 1990) and visual (Neiworth & Wright, 1994) categorization tasks. This might suggest that they are operating in a rule-based, single-dimensional way. However, Smith, Redford, Haas, Coutinho, and Couchman (2008) evaluated this possibility and found that while humans approach same-different tasks in a rule-based way, macaques used a similarity-based strategy. Humans distinguished 0-disparity pairs from pairs that had any discernable disparity, indicating that they were using a same-different rule. Monkeys distinguished similar (including 0-disparity and low disparity) pairs from high-disparity pairs, indicating that they could attend to the sameness-difference dimension but were not using it as a criterial rule.

It is possible that the present tasks actually point not just to preferential performance patterns by monkeys, but also to structural constraints on rule use by monkeys. In related research, Smith et al. (2004) gave monkeys category tasks that demanded nonoptionally the use of single-attribute or correlational rules, and they also found that monkeys were impoverished in rule use relative to humans. Likewise, in Smith et al. (in press), monkeys were far slower than humans at appreciating single-dimensional category solutions even when the use of rules was mandatory.

Another approach to fostering rule use was used by Shepard et al. (2001) and by Harlow (1949). Shepard et al. gave humans a series of consecutive category problems of the same type so they could gain skill in this area. Humans came to acquire rules nearly instantaneously. Harlow (1949) had some success with this same approach in fostering the acquisition or learning sets or rules by monkeys. Our preliminary use of this technique in Experiment 3 gave no hint that the monkeys were developing learning sets for the use of criterial rules, though it did show that the repeated tasks sharpened monkeys' general accuracy. At least to the point of the ten task repetitions used here, Harlow's successful fostering of rule use did not extend for monkeys into the domain of categorization.

It is interesting to consider from a cognitive and neuroscience perspective the differences between humans and monkeys that created the present cross-species difference. One possibility is that humans' language resource, and their verbalization of hypotheses selected and rules applied, lies behind their present emphasis on CA strategies in the present tasks. In fact, the verbal rules for CA solutions are trivially simple. Shepard et al. (1961) raised the same possibility. Feldman (2000) also supported the language hypothesis by showing that the rank

order of difficulty of category tasks for humans was closely tied to the complexity of the verbal rule that described the task's solution. Interestingly, Smith et al. (2004) showed that this language-based rank order of difficulty did not hold for monkeys, sharpening interest in a language hypothesis.

A related or higher-level description of the species difference comes from current theory in the cognitive neuroscience of category learning (Ashby & Ell, 2000; Ashby et al., 1998). Humans may have a more highly developed declarative and intentional category-learning system because the frontal cortical structures (e.g., prefrontal cortex and the anterior cingulate gyrus) that mediate that system (by managing the processes of hypothesis memory, testing, and switching) are highly developed. In contrast, monkeys—with small frontal cortices and less highly developed executive-attention abilities—might have a less robust and more rudimentary intentional category-learning system.

There is one theoretical reservation about this conclusion. Humans—especially humans learning intentionally—have encountered a lifetime of instruction and bias toward rule-based solutions to the grouping and categorization problems posed by life and the educational system. Humans probably have highly trained and entrained learning sets to test hypotheses and to find and adopt rule-based strategies, especially under the clear instructions that characterized the intentional learning condition here. Monkeys surely lack that lifetime experience, and moreover one cannot instruct monkeys as directly as one can instruct humans. Possibly, therefore, we are seeing in the performance of humans this cultural-educational training in operation. (One could explain the developmental progression toward CA responding using this same mechanism and invoking the advent of school at about age 5 for children). It is an interesting chicken-egg problem: do humans appreciate dimensional rules and solutions strongly because of their cultural-educational systems, or do cultural-educational systems appreciate rules because of humans' cognitive makeup? Perhaps the structural features of the rule-based categorization system in the brain – which probably predate educational systems— break the tie in favor of an interpretation based in humans' cognitive makeup.

This discussion raises an interesting issue regarding the evolution of systems of categorization in the primates. The comparative results suggest that there may have been a categorization system, mediated by processes within the striatum or basal ganglion, that was shared broadly across the primates and from which the human categorization system arose. Human extended this system by amplification, through a highly developed capacity for sharper, focused attention in category tasks using hypothesis testing and rule-based solutions.

We conclude by pointing out that there are many substantial benefits to this additional capacity once developed by humans. It could aid in the memory of categories or problem solutions from situation to situation, by coding them as simply as possible. This could improve the performance of ecological category tasks and support the efficiency and safety of foraging. Notice that even in the present, non-ecological studies here, rule-using humans were highly accurate, but FR-responding humans (incidental learners) and FR-responding monkeys were less accurate. It should also be noted that

The benefits of CA strategies and rule use could have a critical communicative and cultural role. A human CA processer could teach conspecifics or offspring a simple, single-dimensional, explicit rule – for example, "leaves of three, let it be" to explain the avoidance behavior for poison ivy. A hypothesis that is well-tested by generations of humans, or even by one human, can be directly communicated and is therefore immediately available to others. Cultural transmission and cultural advances can be substantially accelerated in this way. In contrast, an FR categorizer might have to communicate a category's intension by presenting many exemplars or by attempting a probabilistic featural description that was far less

parsimonious. In a sense, a category understood by family-resemblance relationships is more private and perceptual, useful to a single individual but less communicable. Indeed, humans often have great difficulty expressing why they perform as they do when they perform complex FR tasks using dot distortions, artificial grammars, and even multi-dimensional tasks such as those used here. Therefore, categories of this kind might tend to slow cultural transmission and advancement. In fact, it is well known that apes and monkeys, compared to humans, communicate far less about rules, roles, and functions in tool-use and problem-solving situations, and this is often used to explain why the culture of nonhuman primates displays such profound stasis. Each individual has to relearn, on its own, the cognitive knowledge toolkit of the species, whereas humans can transmit this cognitive toolkit directly and quickly. Thus, one sees that the human evolution toward the confluence of rules, communication, and language likely had multiple benefits that are well illustrated from the perspective of categorization.

Acknowledgments

References

Ahn W, Medin DL. A two-stage model of category construction. Cognitive Science 1992;16:81–121.

Ashby FG, Alfonso-Reese L, Turken U, Waldron E. A neuropsychological theory of multiple systems in category learning. Psychological Review 1998;105:442–481. [PubMed: 9697427]

Ashby FG, Ell S. The neurobiology of human category learning. Trends in Cognitive Science 2001;5:204–210.

Ashby FG, Maddox WT. Complex decision rules in categorization: Contrasting novice and experienced performance. Journal of Experimental Psychology: Human Perception and Performance 1992;18:50–71.

Ashby FG, O'Brien J. Category learning and multiple memory systems. Trends in Cognitive Sciences 2005;9:83–89. [PubMed: 15668101]

Ashby FG, Queller S, Berretty P. On the dominance of unidimensional rules in unsupervised learning. Perception and Psychophysics 1999;61:1178–1199. [PubMed: 10497436]

Ashby FG, Spiering B. The neurobiology of category learning. Behavioral and Cognitive Neuroscience Reviews 2004;3:101–113. [PubMed: 15537987]

Brooks, LR. Nonanalytic concept formation and memory for instances. In: Rosch, E.; Lloyd, BB., editors. Cognition and categorization. Hillsdale, NJ: Erlbaum; 1978. p. 169-211.

Cerella J. Visual classes and natural categories in the pigeon. Journal of Experimental Psychology: Human Perception and Performance 1979;5:68–77. [PubMed: 528931]

Cook RG, Smith JD. Stages of abstraction and exemplar memorization in pigeons' category learning. Psychological Science 2006;17:1059–1067. [PubMed: 17201788]

Erickson MA, Kruschke JK. Rules and exemplars in category learning. Journal of Experimental Psychology: General 1998;127:107–140. [PubMed: 9622910]

Feldman J. Minimization of Boolean complexity in human concept learning. Nature 2000;407:630–633. [PubMed: 11034211]

Harlow HF. The formation of learning sets. Psychological Review 1949;56:51–65. [PubMed: 18124807]

Herrnstein RJ, Loveland DH, Cable C. Natural concepts in pigeons. Journal of Experimental Psychology: Animal Behavior Processes 1976;2:285–301. [PubMed: 978139]

Huber, L. Visual categorization in pigeons. In: Cook, R., editor. Avian visual cognition. Medford, MA: Robert Cook and Comparative Cognition Press; 2001. [On-line]. www.pigeon.psy.tufts.edu/avc/huber/

Huber L, Lenz R. A test of the linear feature model of polymorphous concept discrimination with pigeons. Quarterly Journal of Experimental Psychology 1993;46B:1–18.

Jitsumori M. Category discrimination of artificial polymorphous stimuli based on feature learning. Journal of Experimental Psychology: Animal Behavior Processes 1993;19:244–54.

Jitsumori, M. Discrimination of artificial polymorphous categories in humans and nonhumans. In: Hayes, SC.; Hayes, LJ., editors. Behavior analysis of language and cognition. 1994. p. 91-106.

Jitsumori M. Prototype effect and categorization of artificial polymorphous stimuli in pigeons. Journal of Experimental Psychology: Animal Behavior Processes 1996;22:405–418. [PubMed: 8865609]

Kemler DG. Classification in young and retarded children's discriminative learning: The primacy of overall similarity relations. Child Development 1982;53:768–779. [PubMed: 7094679]

Kemler Nelson, DG. Holistic and analytic modes in perceptual and cognitive development. In: Tighe, T.; Shepp, BE., editors. Perception, Cognition, and development: Interactional analyses. Hillsdale, N.J: Erlbaum; 1983a. p. 1-37.

Kemler Nelson DG. Exploring and re-exploring issues of integrality, Perceptual sensitivity, and dimensional salience. Journal of Experimental Child Psychology 1983b;36:365–379.

Kemler Nelson DG. The effect of intention on what concepts are acquired. Journal of Verbal Learning and Verbal Behavior 1984;23:734–759.

Kemler Nelson DG. When category learning is holistic: A reply to Ward and Scott. Memory and Cognition 1988;16:79–84.

Kemler Nelson, DG.; Smith, JD. Analytic and holistic processing in reflection-impulsivity and cognitive development. In: Globerson, T.; Zelniker, T., editors. Cognitive style and cognitive development. Norwood, NJ: Ablex; 1989. p. 116-140.

Krechevsky L. "Hypotheses" in rats. Psychological Review 1932;39:516–532.

Lashley KS. An examination of the "continuity theory" as applied to discriminative learning. Journal of General Psychology 1942;26:241–265.

Lea SEG, Wills AJ. Use of multiple dimensions in learned discriminations. Comparative Cognition and Behavior Reviews 2008;3:115–133.

Love BC. Comparing supervised and unsupervised category learning. Psychonomic Bulletin and Review 2002;9:829–835. [PubMed: 12613690]

Mackintosh NJ. A theory of attention: Variations in the associability of stimuli with reinforcement. Psychological Review 1975;82:276–298.

Medin, DL. A theory of context in discrimination learning. In: Bower, GH., editor. The psychology of learning and motivation. Vol. 9. New York: Academic Press; 1975. p. 269-315.

Murphy, GL. The big book of concepts. Cambridge: MIT Press; 2002.

Neiworth JJ, Wright AA. Monkeys (Macaca mulatta) learn category matching in a nonidentical same-different task. Journal of Experimental Psychology, Animal Behavior Processes 1994;20:429–35. [PubMed: 7964525]

Nosofsky RM, Palmeri TJ, McKinley SK. Rule-plus exception model of classification learning. Psychological Review 1994;101:53–79. [PubMed: 8121960]

Pearce, JM. Discrimination and categorization. In: Mackintosh, NJ., editor. Animal learning and cognition. New York: Academic Press; 1994. p. 109-134.

Regehr G, Brooks LR. Category organization in free classification: The organizing effect of an array of stimuli. Journal of Experimental Psychology: Learning, Memory, and Cognition 1995;21:347–363.

Richardson WK, Washburn DA, Hopkins WD, Savage-Rumbaugh ES, Rumbaugh DM. The NASA/LRC computerized test system. Behavior Research Methods, Instruments, and Computers 1990;22:127–131.

Roberts AC. Comparison of cognitive function in human and nonhuman primates. Cognitive Brain Research 1996;3:319–327. [PubMed: 8806033]

Rosch EH, Mervis CB. Family resemblances: Studies in the internal structure of categories. Cognitive Psychology 1975;7:573–605.

Rosch EH, Mervis CB, Gray WD, Johnson DM, Boyes-Braem P. Basic objects in natural categories. Cognitive Psychology 1976;8:382–439.

Rumbaugh DM, Richardson WK, Washburn DA, Savage-Rumbaugh ES, Hopkins WD. Rhesus monkeys (*Macaca mulatta*), video tasks, and implications for stimulus-response spatial contiguity. Journal of Comparative Psychology 1989;103:32–38. [PubMed: 2924530]

Semendeferi K, Lu A, Schenker N. Humans and great apes share a large frontal cortex. Nature Neuroscience 2002;5:272–276.

Shepp, BE. The analyzability of multidimensional objects: Some constraints on perceived structure and attention. In: Tighe, T.; Shepp, BE., editors. Perception, cognition, and development: Interactional Analysis. Hillsdale, NJ: Erlbaum; 1983.

Shepp, BE.; Burns, B.; McDonough, D. The relation of stimulus structure to perceptual and cognitive development: Further tests of a separability hypothesis. In: Becker, J.; Wilkening, F., editors. The integration of information by children. Hillsdale, NJ: Erlbaum; 1980. p. 113-145.

Smith JD, Beran MJ, Crossley MJ, Boomer JT, Ashby FG. Implicit and explicit category learning by macaques (*Macaca mulatta*) and humans (*Homo sapiens*). Journal of Experimental Psychology: Animal Behavior Processes. In press.

Smith JD, Kemler Nelson DG. Is the more impulsive child a more holistic processor? A reconsideration. Child Development 1988;50:719–727.

Smith JD, Minda JP. Prototypes in the mist: the early epochs of category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition 1998;24:1411–1436.

Smith JD, Minda JP, Washburn DA. Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins (1961) tasks. Journal of Experimental Psychology: General 2004;133:398–414. [PubMed: 15355146]

Smith JD, Shapiro JH. The occurrence of holistic categorization. Journal of Memory and Language 1989;28:386–399.

Smith JD, Tracy J, Murray MJ. Depression and category learning. Journal of Experimental Psychology: General 1993;122:1–16.

Smith JD, Redford JS, Haas SM, Coutinho MVC, Couchman JJ. The comparative psychology of same-different judgments by humans (*Homo sapiens*) and monkeys (*Macaca mulatta*). Journal of Experimental Psychology: Animal Behavior Processes 2008;34:361–74. [PubMed: 18665719]

Smith L. Importance of the overall similarity of objects of adults' and children's classifications. Journal of Experimental Psychology: Human perception and performance 1981;7:811–24. [PubMed: 6457093]

Smith LB, Kemler DG. Developmental trends in free classification: Evidence for a new conceptualization. Journal of Experimental Child Psychology 1977;24:279–298.

Sutherland, NS.; Mackintosh, NJ. Mechanisms of animal discrimination learning. New York: Academic Press; 1971.

von Fersen L, Lea SEG. Category discrimination by pigeons using five polymorphous features. Journal of the Experimental Analysis of Behavior 1990;54:69–84. [PubMed: 16812617]

Waldron EM, Ashby FG. The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. Psychonomic Bulletin & Review 2001;8:168–176. [PubMed: 11340863]

Ward TB. Response tempo and separable-integral responding: Evidence for an integral-to-separable processing sequence in visual perception. Journal of Experimental Psychology: Human Perception and Performance 1983;9:103–112. [PubMed: 6220116]

Ward TB, Scott J. Analytic and holistic modes of learning family-resemblance concepts. Memory and Cognition 1987;15:42–54.

Washburn DA, Rumbaugh DM. Testing primates with joystick-based automated apparatus: lessons from the language research center's computerized test system. Behavior Research Methods, Instruments, & Computers 1991;24:157–164.

Washburn DA. Stroop-like effects for monkeys and humans: Processing speed or strength of association? Psychological Science 1994;5:375–379. [PubMed: 11539143]

Wasserman EA, Kiedinger RE, Bhatt RS. Conceptual behavior in pigeons: categories, subcategories, and pseudocategories. Journal of Experimental Psychology: Animal Behavior Processes 1988;14:235–246.

Wittgenstein, Ludwig. Philosophical investigations. Cambridge UK: Blackwell Publishing; 1953/2001.

Wright AA, Shyan MR, Jitsumori M. Auditory same/different learning by monkeys. Animal Learning & Behavior 1990;18:287–94.

Category A

Category B

0 0 0 0

1 1 1 1

0 1 0 0

1 0 1 1

0 0 1 0

1 1 0 1

0 0 0 1

1 1 1 0

Test

0 1 1 1

1 0 0 0

Figure 1.
The logical structure of the categories used in Experiments 1–3. In the first column (i.e., stimulus dimension) for each category, 0 and 1 are the defining featural values (i.e., the criterial attributes) for Categories A and B, respectively. In columns 2–4, Category A and B members, respectively, predominantly take on featural values of 0 and 1. The test stimuli place in opposition the defining feature of one category with the modal featural values of the other category.
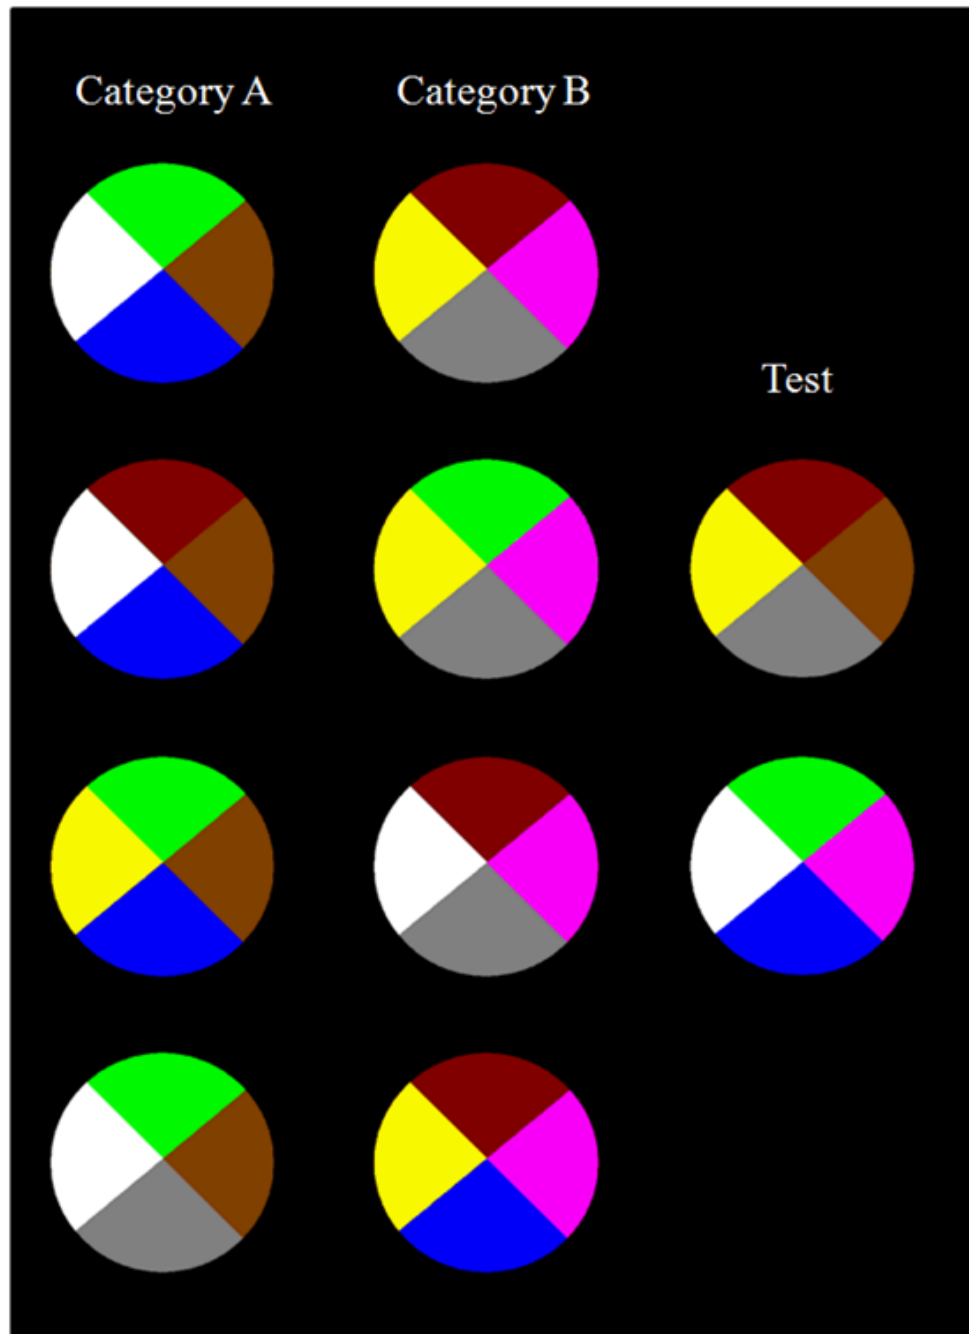
Figure 2.
Illustration of the stimuli used in Experiments 1–3. Each circle is an instantiation of the logical structure of one of the stimuli shown in Figure 1. The circles from top to bottom correspond to the stimuli taken from top to bottom in Figure 1. The first through fourth values of each stimulus in Figure 1, respectively, correspond to the four disk sectors of each circle, beginning with the rightmost sector and proceeding counterclockwise.
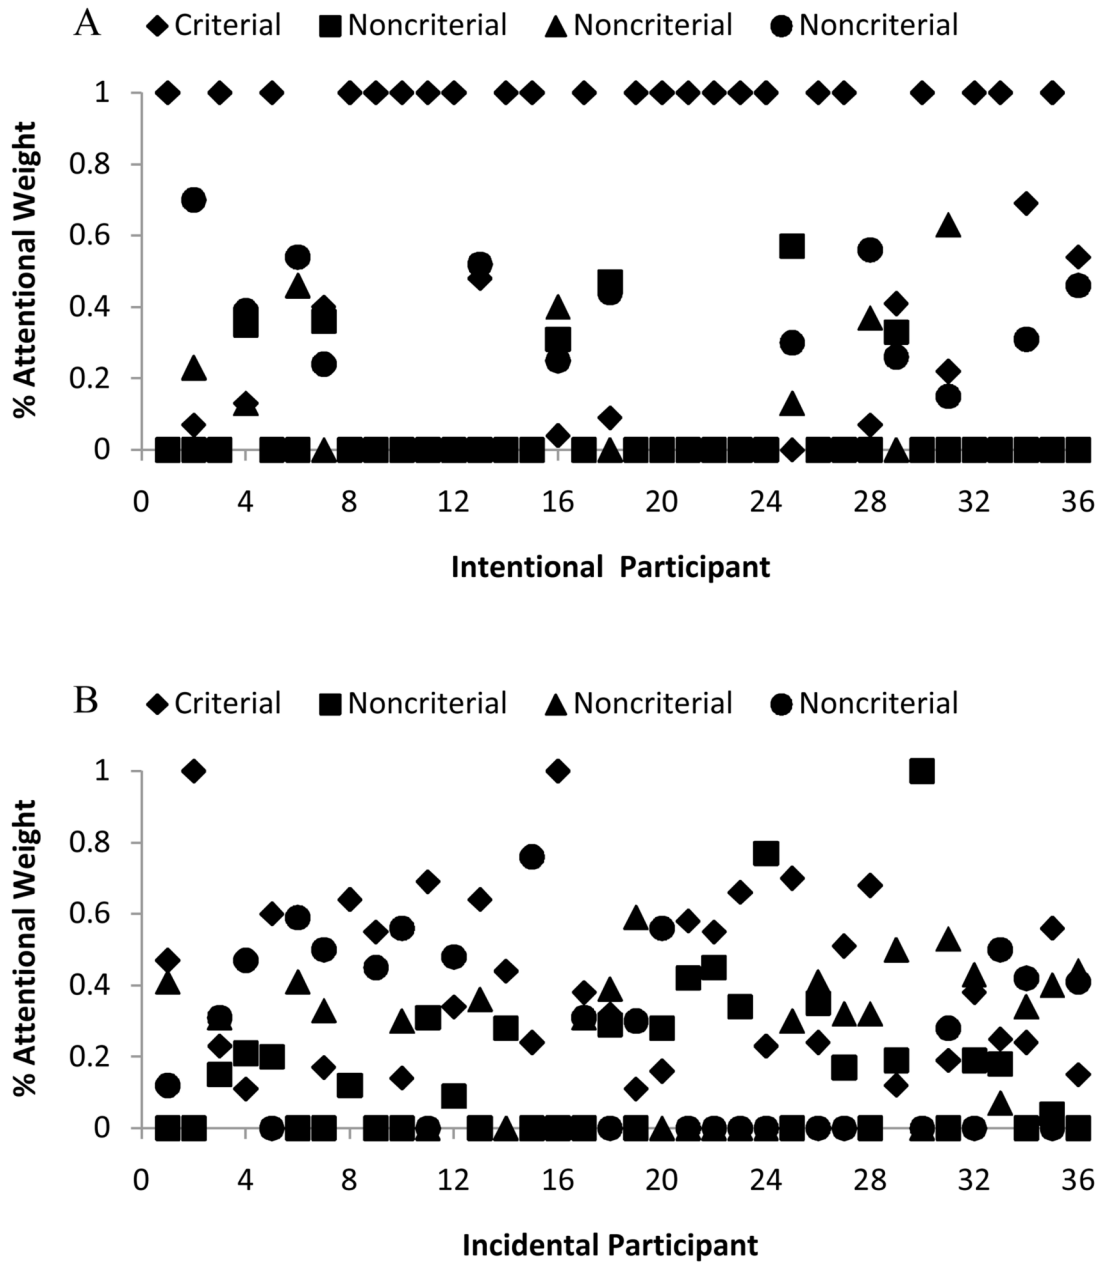
Figure 3.
Results from formal modeling for intentionally learning (A) and incidentally learning (B) humans. Each X-axis value represents one participant, so that each column of points on the graph belongs to one individual. The Y-axis represents the average estimated attentional weight of each participant for each stimulus dimension. The four attentional weights given always sum to 1.0, denoting the presumed total attentional capacity of the participant.
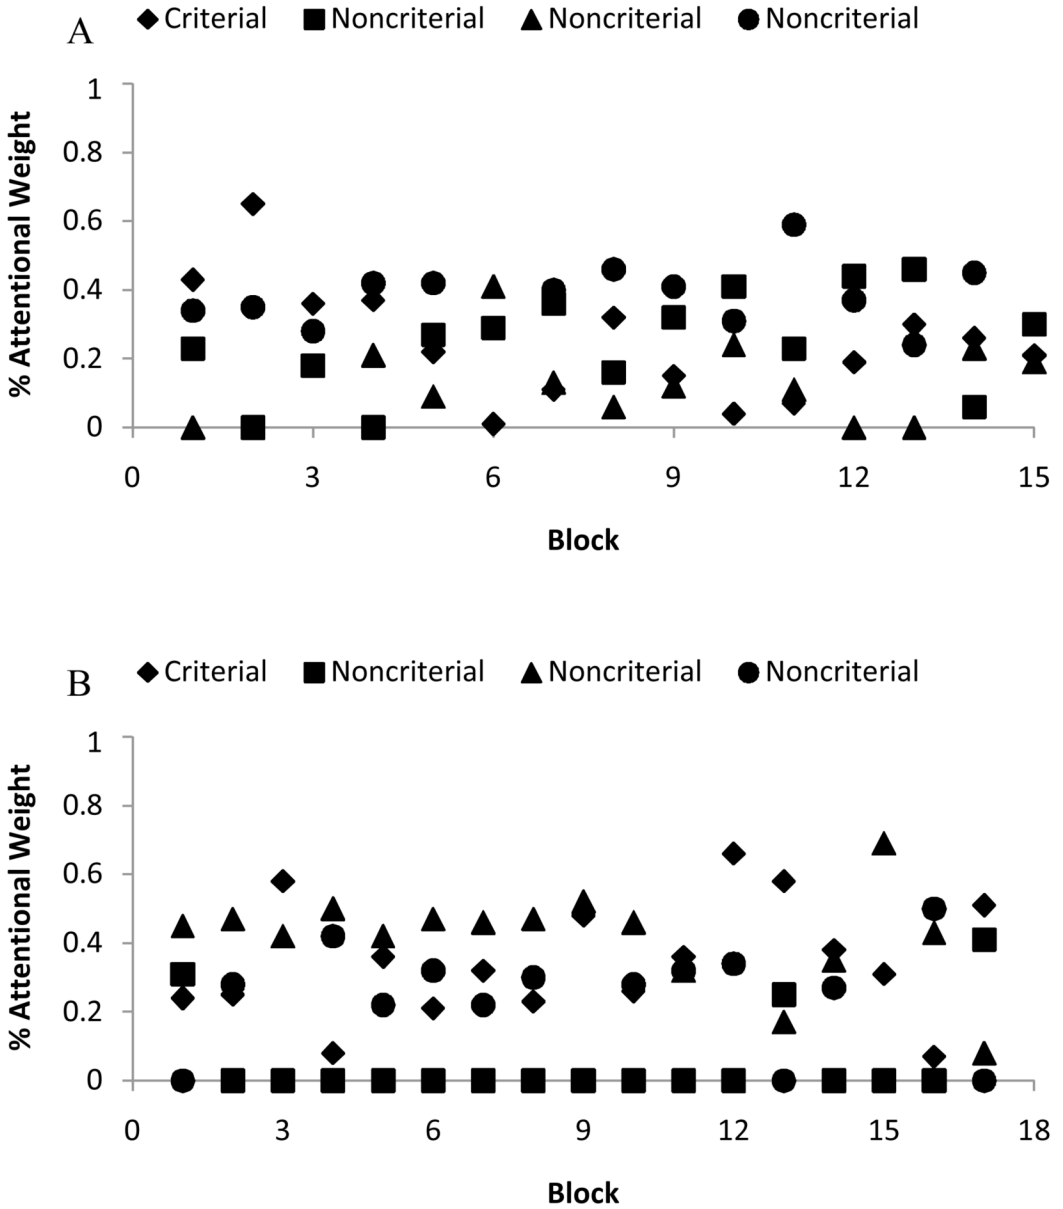
Figure 4.
Results from formal modeling for monkey Murph (A) and Lou (B). Each column of data represents the animal's estimated attentional strategy over 100 trials. The Y-axis represents the average estimated attentional weight for each stimulus dimension. The four attentional weights given always sum to 1.0, denoting the animal's presumed total attentional capacity.
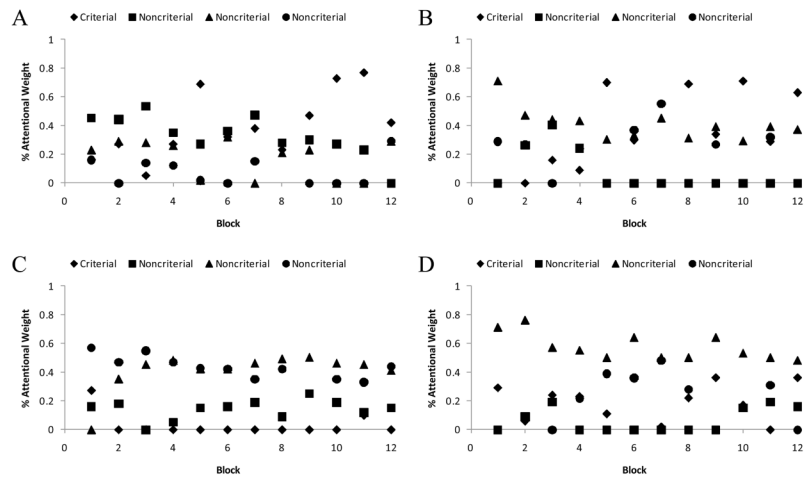
Figure 5.
Results from formal modeling for monkey Murph's first (A) and last (B) sessions, and monkey Lou's first (C) and last (D) sessions, as depicted in Figure 4.