

Diversion of attention in everyday concept learning: Identification in the service of use

LEE R. BROOKS AND ROSEMARY SQUIRE-GRAYDON
McMaster University, Hamilton, Ontario, Canada

AND

TIMOTHY J. WOOD
Medical Council of Canada, Ottawa, Ontario, Canada

Many people tend to believe that natural categories have perfectly predictive defining features. They do not easily accept the family resemblance view that the features characteristic of a category are not individually sufficient to predict the category. However, common category-learning tasks do not produce this *simpler-than-it-is* belief. If there is no simple classification principle in a task, the participants know that fact and can report it. We argue that most category-learning tasks in which family resemblance categories are used fail to produce the everyday simpler-than-it-is belief because they encourage analysis of identification criteria during training. To simulate the learning occurring in many natural circumstances, we developed a procedure in which participants' analytic activity is diverted from the way in which the stimuli are identified to the use to which the stimuli will be put. Finally, we discuss the prevalence of this *diverted analysis* in everyday categorization.

Anyone who teaches introductory cognition knows that many students do not readily accept the complexity of natural concepts. When confronted with the question, "How do you know an animal is a dog?" they respond with such statements as "because it barks" or "because it wags its tail." The instructor can easily show that although these attributes may be characteristic, they are neither necessary nor sufficient for category membership. However, it is much harder to persuade students that they could not have come up with a defining feature if they had taken time to work on it. In short, it takes persuasion to establish the grounds for such ideas as prototypes or family resemblances. Students obviously are fully capable of identifying members of natural categories but have decidedly oversimplified ideas about their criteria for doing so.

Although this *simpler-than-it-is* belief is evident in everyday situations, we do not see it in common category-learning tasks in the laboratory. In these induction tasks, if there is not a simple classification principle, the participants know that fact and can report it. In fact, the persistence with which participants seek simple rules in these tasks, despite their rarity for natural categories, is a problem that calls for an explanation. Starting with the seminal work of Rosch (e.g., Rosch, 1975) on prototypes and the subsequent introduction of instances as bases for categorization (e.g., Brooks, 1978; Medin & Schaffer, 1978), arguments have been made that identifying the members of everyday categories requires something other than simple rules. This position is an echo of strong

arguments from other fields that simple rules are unlikely to be a basis for the learning of everyday concepts (e.g., Putnam, 1973; Wittgenstein, 1953). Despite the force of these arguments for the academic community, subsequent evidence has demonstrated that laboratory participants are still strongly attached to the idea that categories have defining features (e.g., Ahn & Medin, 1992; Medin, Wattenmaker, & Hampson, 1987; Regehr & Brooks, 1995; reviewed in Murphy, 2002, pp. 126–134).

In this article, our first general question is how people get as far as the university classroom believing that complex natural categories are susceptible to easy definition. That is, how can people function in a complex world and yet think it is simple? It cannot be just a matter of people wishing to think of the world as a simple place. If people sought simple rules in the everyday world with anything close to the persistence that they show in laboratory tasks, they would know that simple rules do not fit. Our second general question is why our usual laboratory tasks do not produce the simpler-than-it-is belief. What is it that is different about laboratory tasks that induces more analysis, leading to the discovery that single features are not perfect predictors?

The general answer that we will give to the first question is that people do not know about the complexity of natural categories because they do not explicitly try to solve the problem of how to determine category membership. Instead, people generally think about some aspect of dealing with an object that is subsequent to actually

identifying it. The subsequent activity that will be the topic of this article is a *focus on use*: An identification decision can be effectively backgrounded by attention to the use to which an object will be put. B. H. Ross (1996, 1999) has shown that subsequent use can change what is known about a category and can provide different bases of categorization. Minda and Ross (2004) have introduced a procedure, indirect category learning, in which they show that when attention to use has backgrounded classification itself, people accomplish identification with a less explicit procedure, such as generalizing from previous instances or relying on a set of merely characteristic features. Relevant to the concerns of this article is what happens when people who have paid little attention to the basis of the categorization of naturally occurring objects are asked an explicit question about identification, such as “how do you know an animal is a dog?” We intend to show that they have to develop an explicit answer, probably for the first time, and that this answer often is only loosely constrained by the criteria they actually used in their previous identifications.

Among the many possible influences on what people say about the identification of categories is the impression of *discriminative similarity* gained from having dealt with past instances—that is, the impression of coherence that comes from finding similarity among the items within a category and finding contrast to items in other categories. It would be easy to believe that this similarity stems from features that are completely consistent. This impression could be gained during initial exposure to the stimuli or from reflecting on them at the moment of test. Another possible influence on what people say about categories is a prior belief that unless it is proven otherwise, all categories have common elements, a belief called *psychological essentialism* by Medin and Ortony (1989; see also Kalish, 2002, for a discussion of the sense in which this belief is held). However, regardless of whether its source is an impression of discriminative similarity or a judgment bias, when explicitly challenged to define a category, students have not done enough analysis to disconfirm the sufficiency of individual cues. As a result, they can express high confidence that there is some common element and, when pressed, propose a simple hypothesis that could easily have been eliminated if it had been considered in the presence of the stimuli. That is, they start as if they are trying to solve the problem for the first time.

Our answer to the second general question is that most laboratory tasks do not result in the usual everyday *failure to analyze* because they strongly encourage participants to analyze their own identification procedures during training. Participants are put in a formal procedure that suggests they are under an analytic microscope; they have the existence of new categories signaled to them and, perhaps most critically, are given nothing else to think about than identifying the new categories. We believe that these conditions are likely to produce considerable analysis by the participants regarding the basis of identification. However, regardless of whether the absence of attempts to analyze under everyday conditions or the analytic press of

laboratory tasks is the critical factor, we have an obvious discrepancy between the behavior of people in standard laboratory tasks and their behavior toward natural categories. Without controlling for the discrepancy between analytic effort in the standard classification-learning tasks and that in many everyday conceptual acquisition situations, the relation of laboratory work to everyday behavior will be unclear.

This part of our work joins a movement in the category-learning literature to capture a greater variety of everyday learning conditions. Examples include the learning-by-inference task of Yamauchi and Markman (1998, 2000), reviewed by Markman and Ross (2003), and the analytic/integrative contrast reported by Whittlesea, Brooks, and Westcott (1994). In each of these investigations, as in the present study, what was learned about a family resemblance structure depended on the learning activity. In the learning-by-inference task in Yamauchi and Markman (1998) and Markman and Ross (2003), the learners predicted the missing value in a set of items and showed later evidence of having gained a prototype-like knowledge of the category. The learners in Whittlesea et al. engaged in both an integrative activity with the learning items and an activity in which they evaluated the categorical evidence provided by individual features in the same learning items. That is, they engaged in both integrative and analytic activities with the same items. Later, they showed transfer that was mainly dependent on similarity to old items or that was mainly dependent on similarity to an implicit prototype, depending on whether integrative or analytic coding was cued by the details of the transfer task. In the present study, diverting attention to the use to which an item would be put changed the learner’s metacognition—that is, the explicit knowledge of simple facts about the structure of the items just learned.

The Diverted Analysis Procedure

We propose that in order to simulate in the laboratory the *solving it for the first time when explicitly asked* performance, we need to meet two conditions.

1. In the study phase, while our participants are becoming familiar with the stimuli, we must *divert their analytic activity* onto some aspect of dealing with the stimuli other than categorical identification. With full attention to the identification task in an experiment, participants display a persistent effort to discover explicit identification rules that is quite unlike their usual dealings with everyday categories.

2. We must ensure *supportive learning conditions* during the initial acquisition phase. In the world, people are usually free to attend to an identification problem if it gives them difficulty, but they do attend to it only if they start making mistakes or if they encounter stimuli for which a categorization is not readily apparent. In many situations in the world, there is considerable contextual support for identification, such as naming of a new object by other people, until familiarity is gained with at least some category members. Furthermore, there are often redundant cues available for distinguishing new categories, particu-

larly feature instantiations that are unique to particular categories. The present experiments are designed to simulate this *supportive stimulus structure* in the laboratory.

As an illustration, consider the task and materials shown in Figure 1. In this task, participants are told that there are two categories of animals: bleebes and ramuses. The bleebes can move only on diagonals (like a bishop in chess), and the ramuses can move only on columns and rows (like a rook in chess). The task is to determine which of the two pieces will take the fewest moves to reach the finish, labeled with an F, without entering any of the squares marked out with an X. In each subsequent trial, a slightly different board is used: different position of the finish, different obstacles, different starting positions of the animals, and most important, different animals. The full set of animals is shown in Figure 2 (participants never see this array of all eight animals at the same time). Gradually, during subsequent passes through the training stimuli, the presentation is shifted to the animals' being unlabeled and occurring one at a time, rather than in pairs, as shown in Figure 1. Overall, the training procedure is intended to emphasize the use to which the animals are put and to give enough initial support for identification to trivialize the problem of identification.

In the subsequent test phase, we have three potential indications of a lack of analysis during training—that is, three indications that the participants are starting explicit analysis of identification criteria for the first time when challenged at the beginning of the test phase.

1. In response to the question, “Do you think that there is at least one feature or physical characteristic that all bleebes have in common that no ramuses have?” participants with a general, unexamined impression of high discriminative similarity or prior bias toward assuming commonality will be likely to say “yes.” However, the correct answer for the *obvious* dimensions in these animals, given in the Appendix, is that no feature completely predicts the categories (although, as will be discussed later, there are some nonsalient predictive dimensions that can be discovered with enough analysis).

2. In response to the request, “Please give a rule that would most accurately separate bleebes from ramuses,” such participants will give a single-feature rule that can easily be eliminated if considered in the presence of the stimuli.

3. In response to the question, “Did you try to figure out a rule that separated bleebes from ramuses?” the participants will say that little or no time was spent trying to discover a rule during the chessboard task.

EXPERIMENT 1

Task Control of Analysis During Initial Training

In our first experiment, we contrasted the performance of a *diverted analysis* group with that of groups run in two common laboratory methods: an *explicit induction* group with instructions to look for a rule and a group with instructions to *memorize* the training items. In both of these procedures, the participants' attention is allowed to focus on the identification problem itself, rather than

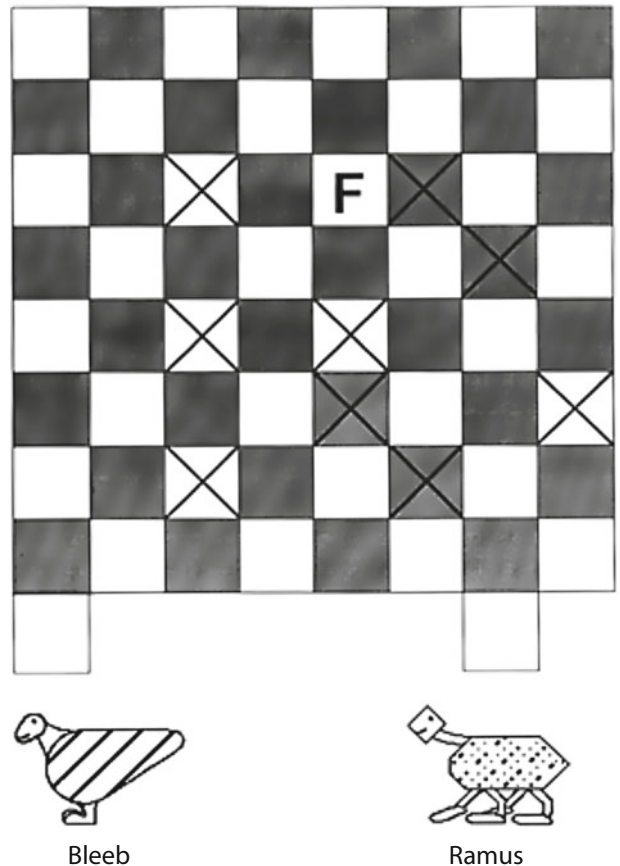


Figure 1. An example of the materials used to divert attention from the act of identification to the use to which the exemplars would be put.

being diverted elsewhere. As a result, they will be more likely during the acquisition phase to identify any common elements or cite their failure to find any common elements, despite having tried. In neither case are they likely to propose the sufficiency of a simple but incorrect rule and inappropriately express high confidence that there is a single discriminative element.

Our interest in the explicit induction group was to provide an extreme contrast against which to evaluate the performance of the diverted analysis group. The explicit induction procedure was as close a match as possible to that used for the diverted analysis group, with the exceptions that they were given instructions to look for a rule and were not provided with the moves or the boards as a diverting task.

Our interest in the memorize group was to evaluate whether instructions to memorize the categorization of items might induce the participants to bypass analysis of the featural basis of classification. Memorization is a procedure that has been suggested to induce a less analytic approach to learning (e.g., Brooks, 1978; Higham, 1997; Kemler-Nelson, 1989; Reber, 1989; Ward & Scott, 1987). Because we were interested in the conditions that might induce learning of the categories without explicit knowledge of the basis of categorization, memorization




Figure 2. The local style stimuli used in Experiment 1. None of the obvious dimensions (two/four legs, round/angular bodies, dotted/striped markings, round/angular heads) were perfectly correlated with the categories.

was obviously a relevant condition to evaluate. Memorization of the appearance of the surface features of items might indeed eliminate analysis for single-category classification procedures, such as artificial grammar learning, because there is nothing in such tasks that suggests the presence of a category. However, when several items are labeled with one explicit response and other items labeled with another, as happened in the present experiment, the participants might easily (and generally correctly) presume that categories of items are important in the experiment. These signals, together with the freedom to deploy full attention on the identification task, might be all that is necessary to induce university students in a laboratory setting to attempt to discover the basis of categorization. Possibly, such conditions might not be sufficient to induce children to attempt to discover rules, rather than directly comply with the instruction to memorize. However, for the highly verbal undergraduate participants in our experiments, more than just instructions to memorize might be required to prevent the participants from trying to explicitly analyze the basis of categorization during learning. Greater signs of analysis during learning in the memorize group than in the diverted analysis group would support this possibility.

To ensure easy identification during the initial stages of acquisition, we provided what we referred to as a *supportive stimulus structure*, shown in Figure 2. A critical aspect of this set is that the specific manifestation of a feature, its instantiation, is not the same in both categories. We believe that this is an extremely common and crucial characteristic of most everyday, concrete categories (Brooks & Hannah, 2006; Hannah & Brooks, 2006a, 2006b; a very similar contrast is that between global and local properties in representation, as shown in property verification experiments by Solomon & Barsalou, 2001, 2004). Both birds and humans have two legs, but clearly, they are not the same two legs. All we would need to see is the feathered legs of a hawk to know that we are not dealing with a human or even another species of bird. Such strong predictors allow for rapid and reliable identification under degraded viewing conditions, possibly even with a single feature. In contrast, a more general descriptor, such as *two legs*, applies to humans and birds but, in itself, picks out neither of them. This broad scope descriptor supports generalization to new or atypical items and allows the discovery of higher order relations that are critical to understanding the structure of a domain or other abstract principles. *Instantiated features*, the perceptually specific manifestations, are good for identifying a bird or even a specific bird; *informational features*, the more abstract descriptors, are useful for learning what types of movement a bird can make, as well as for communications and thinking about the whole category. Although the perceptual specificity of an instantiated feature is useful under many diagnostic conditions, this specificity means that it is a part of only a limited number of items within the category (i.e., low categorical validity). To cover a whole category, many instantiations would be necessary, which makes for difficulties in description, communication, and instruction. Here, the generality of informational features becomes an advantage, despite the fact that they may also apply to a few members from other categories.

In addition to category-specific instantiations, when a characteristic feature in the stimulus set shown in Figure 2, such as two legs, crosses over to the other category, it adopts the *local style* by taking on the appearance of the other values of legs (e.g., four legs) that occur in that category. For the two-legs example, the length of the legs and the appearance of the feet change to those manifested by four legs in that same category. This guarantees that identical surface elements will not occur in two different categories. However, it also means that if the learners redescribed the *two-leg* feature to the less salient *shape of feet*, they would have a completely consistent feature by which to distinguish the two categories. The same argument holds for each of the other dimensions: round/angular body shape versus front-back symmetrical/asymmetrical body shape, stripes/dots versus coarse/fine markings, and rounded/angular versus regular/elongated heads (the informational structure of the high salience dimensions is given in part A of the Appendix). We originally adopted this *perfectly predictive, low-salience dimension* structure because we believed that a greater appearance of similarity within a category (e.g., all of the legs within a category

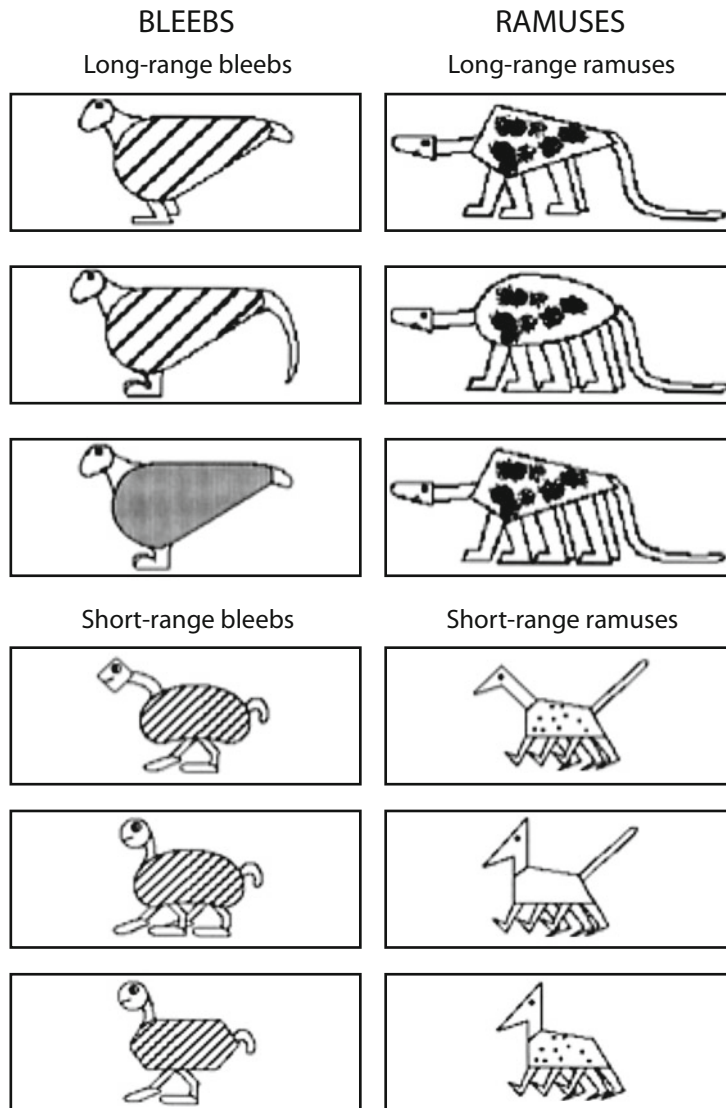


Figure 3. The breeds-within-species stimuli used in Experiment 2. Regardless of its being a short- or a long-range breed, members of the bleeb species usually have two legs, a rounded body, a rounded head, stripes, and a short tail.

were very similar to one another, rather than having the multiple forms shown in Figure 3) would be required to avoid analytic behavior than is afforded by many standard stimulus structures. In fact, the basic findings of Experiment 1 can also be demonstrated with much more varied sets of stimuli (e.g., those in Figure 3, used in Experiment 2). However, we decided to report the data using the present local style stimulus set because it gave us the possibility of one more converging line of evidence regarding analysis by the participants during the acquisition phase. If the low-salience dimensions could be discovered only by analyzing the stimuli during acquisition, mention of these dimensions should occur only in conditions that were likely to encourage such analysis. In fact, only 8 of the 72 participants in this experiment discovered any

of the low-salience dimensions, and they were all in the memorize and explicit induction groups.

Method

Participants. The participants were 72 McMaster students, 24 in each of three groups, taking an introductory psychology course and participating for course credit. The participants were randomly assigned to the three conditions in this experiment and were run individually.

Materials. The stimuli were the drawings of imaginary animals shown in Figure 2, with the informational structure shown in part A of the Appendix. In diverted analysis training, the animals occurred in pairs consisting of one bleeb and one ramus. For the first pass through the five pairs of animals in a given condition, the animals were labeled as shown in Figure 1, with the bleeb always on the left. In the second pass, the animals were still labeled but could occur on either side. For the third and fourth passes, the animals were unlabeled.

beled and, again, could occur on either side. For each of these four passes, the pairs in which the animals occurred were randomly re-matched, with the restrictions that one bleeb and one ramus occurred in each pair and all the animals were used in a single pass of five trials. For the final (fifth) pass, each animal was presented individually, meaning that there were 10 presentations for this pass. The task for these single-item trials was to say how many moves were required for the animal to reach the F. This last pass allowed us to show whether or not the participants could (implicitly) categorize the animals individually, rather than only in contrast to another animal.

For the diverted analysis group, each presentation card also contained an illustration of a modified chessboard, the same size and form as those shown in Figure 1. The first two presentations contained no obstacles (squares with an X); all the subsequent presentations contained from 2 to 14 obstacles, with the same presentation order being used for all the participants. The position of the finish (F) square was varied. The intent of this sequence was to provide a few easy trials, with subsequent trials providing noticeably different and more challenging problems. The problems given by the boards were designed so that both animals could reach the F square (this required that the *starting squares* appended to the bottom of the board were of the same color as the F square; for half the trials, these squares were black, and for the other half, they were white), that one animal could always reach the finish in fewer moves than the other, and that the number of moves required would be different for each animal both if properly identified and if misidentified. Critically, these constraints allowed us to determine the category chosen by the learner without directly asking—that is, to determine the learners' categorization while maintaining their attention on the moves, rather than on the basis of response.

In memorize and explicit induction training, each presentation consisted of only the animals, without the chessboard. The pairing, number of passes, and sequence of presentation for these groups was identical to those for the diverted analysis group.

Procedure. In the training phase for the diverted analysis group, a participant was given a response sheet and was told, "You will be working with two kinds of creatures, bleebes and ramuses. Bleebes can only move in diagonal lines, and only on one color, either black or white. Ramuses can only move in straight lines: forward, backward and side to side, and can move on both black and white." The movement pattern was then demonstrated using the first board. "Your job is to decide who, the bleeb or the ramus, gets to the finish square first." The finish square, marked F, on the first board was indicated. "By first, we mean in the *least* number of moves. A move can be any number of squares in a single direction. Changing direction signifies the end of a move." Four moves each for the bleeb and the ramus were then demonstrated. "At times you will encounter blocks on the board. You cannot go through them; you must go around them." The first board with blocks, denoted by Xs on the board, was shown, and the blocks were indicated. "Please write the winner, R for ramus, B for bleeb, and the number of moves for the winner on your sheet. When you are finished with one board, I will give you another. You may use the end of your pencil/pen as a guide, but please do not mark on the boards." The participant's answers, including spontaneous comments, were noted. If the participant entered a wrong answer, he or she was stopped and shown the correct path. For the single-animal presentations (the last 10 presentations), the participant was asked to complete the maze task as quickly as possible, still being accurate, and was told that he or she would be timed (despite brandishing a stopwatch, the experimenter actually recorded other aspects of the participant's response and general behavior, rather than response time). For these final 10 trials, feedback was given; as with the feedback throughout, if the category was incorrect, the participant was told the correct label and was asked to rework the chessboard. The training phase was completed when all 30 presentations had been run. The boards and scoring sheet were taken by the experimenter before the test was begun.

In the memorize training phase, the participant was told, "You will be working with two kinds of creatures, five bleebes and five

ramuses. Each of them will be shown several times, and your job is to memorize which ones are bleebes and which ones are ramuses." For the first pass through the set of animals, the experimenter named each animal (bleeb or ramus) as it was shown and asked the participant to memorize it. When the participant indicated that the animal had been memorized, the card was removed. The procedure was then repeated with the next card, until all five had been shown. In subsequent passes, the animals were randomized in the same order as that used for the diverted analysis group, and the participant was asked to name each animal as it was shown. The participant was given immediate feedback and was asked to memorize the animal. When the participant indicated that the animal had been memorized, the card was placed face down. The procedure was then repeated with the next card, until all the cards had been shown.

In the explicit induction training phase, the presentation order was identical to that for the memorize group (four rounds of paired presentations plus one of single-item presentations). However, the participant was instructed to "please try to find a rule for telling the animals labeled bleebes from those labeled ramuses so that you can correctly identify all the animals in these two categories." As with the memorize group, the participant was given immediate feedback on each trial and the same number of presentations, and the trials were self-paced.

In the test phase, the following questions were asked of all the groups (the phrases in parentheses were said only to the diverted analysis group).

1. a. Try to answer the following question initially with just a Yes or No. (Apart from the way they move,) do you think that there is at least one feature or physical characteristic that all bleebes have exactly in common that no ramuses have?
 - b. So, would you say that the likelihood that there is a single feature common to all bleebes and no ramuses is high or low? What is the likelihood, on a scale of 1–100 that there is a feature common to all bleebes and to no ramuses, where 100% means you're certain there is one feature common to all bleebes, 0% means you're certain there is no feature that all bleebes have in common but no ramus has, and 50% means that you're completely uncertain, you really don't know?
 - c. What was that feature? [To ensure there was no misinterpretation, the experimenter restated the participant's likelihood response—for example, "so you think it is very likely that all bleebes have a rounded body shape?" If the participant said "not all," the experimenter restated the question and recorded this second response.]
 - d. Let's say that you had to tell someone who had never seen bleebes and ramuses how to tell a bleeb from a ramus by looking at it. What rule would you give?
2. Were you working on this rule previously (while you were doing the boards), or did you come up with it now?
3. Did you expect to be asked about the physical features of bleebes and ramuses?

Results

Table 1 shows the percentages of errors on the final 10 trials, the trials on which single animals were presented. These data are intended to show how well our participants performed in correctly identifying the animals just before being asked about any commonalities the categories might have. If, during these trials, the participants had used a single-feature rule based on the high-salience dimensions, they would have made 20% errors. Only 1 participant in any group in this experiment made as many

Table 1
Percentages of Errors on Moves or Identifications During the Final, Single-Animal Trials in Experiments 1 and 2

Experiment	Percentage of Errors
1A. Diverted analysis–local style	2.9
1B. Memory–local style	1.3
1C. Analyze–local style	0.4
2. Diverted analysis–breeds within species	2.4

Note—For Experiment 1, there were 10 single-item trials; for Experiment 2, in which a greater number of stimuli were used, there were 12 such trials. In Experiment 1, if the participants had been following a single-feature rule during this phase, they would have produced 20% errors; for Experiment 2, 16.7%.

as two errors during the 10 single-item test trials, and no participant made both errors that would be expected on the basis of the single-feature rule that they subsequently mentioned. In fact, the accuracy of identification was very high, and there was no evidence that the participants in any group were following a single-feature rule during the last 10 trials.

Table 2 shows the number and percentage of participants answering “yes” to the commonality question: “Apart from the way they move, do you think that there is at least one feature or physical characteristic that all blebs have exactly in common that no ramuses have?” For an answer to be recorded as a yes, the participant also had to have given a rule naming a single feature as sufficient, without an exception mentioned in it, or to have stated that they were sure that there must be such a rule but that they did not know what it was (which no one did in any experiment reported in this article). A small number answered “yes” and named one of the low-salience dimensions that are perfectly correlated with the categories. By a conservative analysis, these responses should all be counted as a yes. By a more selective analysis, these yeses should be dropped, because the intent of asking the question was to discover which participants believed a dimension to be perfectly predictive of the category when it was not (thereby suggesting that they had not been analyzing during acquisition). In Table 2, the selective count of yeses is reported without parentheses, and the more conservative count is reported in parentheses. Analyses were done according to both criteria, and in no comparison in this article did it change a conclusion at a .05 significance level. Unless otherwise mentioned, the χ^2 values reported are for the more conservative counts (and in all cases, the lower χ^2 value). Descriptively, we will use the selective counts, since they are more representative of the purpose for asking the question. The analyses in all cases were χ^2 for independence within a 2×2 table.

The diverted analysis group answered “yes” to the commonality question 87% of the time, significantly more than either the memorize group (38%, $\chi^2 = 9.375$) or the explicit induction group (16%, $\chi^2 = 14.722$). The latter two groups did not differ from one another significantly ($\chi^2 = 0.784$).

The diverted analysis group answered “yes” (“I was working on a rule previously”) to the “did you analyze?” question 61% of the time (Table 3), significantly less often

than either the memorize group (90%, $\chi^2 = 5.132$) or the explicit induction group (96%, $\chi^2 = 8.178$). The latter two groups did not differ significantly from one another ($\chi^2 = 0.463$).

The diverted analysis group answered “yes” to the anticipation question 27% of the time (Table 4), not significantly less often than the memorize group (50%, $\chi^2 = 2.295$), but significantly less often than the explicit induction group (65%, $\chi^2 = 6.505$). The latter two groups did not differ significantly ($\chi^2 = 1.018$).

No analysis of the answers to the rating task (“What is the likelihood, on a scale of 1–100 that there is a feature common to all blebs and to no ramuses?”) will be presented. The results of these ratings proved to be entirely redundant with the simple *yes/no* version of the question and did not show the hoped-for increase in sensitivity.

Discussion

The diverted analysis training procedure produced the simpler-than-it-is response, described at the beginning of this article. This is shown by the high rate at which the diverted analysis group said “yes,” that there was a perfectly discriminative feature. Since this group made only 2.9% errors on the 10 trials preceding the test phase, we know that they were not following a single-feature rule, since that would have resulted in an error rate of 20%. Evidently, then, they used some other means of identification than the simple rule that they later asserted to be sufficient.

We infer that the diverted analysis group generally was not analyzing the basis of categorization during the acquisition phase, since the single-feature hypotheses that they proposed would have been easily rejected had they been doing so. At the least, we can conclude that during acquisition, they were analyzing less effectively than were the other two groups. This conclusion is corroborated by the diverted analysis group’s giving lower ratings than did the other two groups on both the *analyze* (Table 3) and the *anticipation* (Table 4) questions. It is hard to interpret the absolute level of the analyze question, since we did not try to define, for the participants, what we meant by “working on this rule.” Some participants may have answered on the basis of whether or not they had engaged in explicit

Table 2
Responses to the Commonality Question in Experiments 1 and 2

Experiment	No	Yes	% Yes
1A. Diverted analysis–local style	3	21	87
1B. Memory–local style	13	8 (11)	38 (46)
1C. Analyze–local style	16	3 (8)	16 (33)
2. Diverted analysis–breeds within species	5	19	79

Note—The question was, “Apart from the way they move, do you think that there is at least one feature or physical characteristic that all blebs have exactly in common that no ramuses have?” The numbers in parentheses in the “Yes” column represent the numbers of participants saying “yes,” including the participants who discovered one of the low-salience dimensions that were perfectly correlated with the categories. The numbers in parentheses in the “% Yes” column represent the calculations when the participants who discovered the low-salience dimensions were included.

Table 3
Responses to the Analyze Question for Experiments 1 and 2

Experiment	No	?	Yes	% Yes
1A. Diverted analysis–local style	9	1	14	61
1B. Memory–local style	2	3	19	90
1C. Analyze–local style	1	1	22	96
2. Diverted analysis–breeds within species	14	4	6	30

Note—The question was, “Were you working on this rule while you were doing the boards, or did you come up with it now?” The “?” column records answers from the participants who said that they were not sure. The “% Yes” column records the percentages of “yes” (working on this rule previously) responses among the participants who gave either a “yes” or a “no” response.

hypothesis testing, others on whether or not they had incidentally noticed some consistency. However, regardless of the criterion, fewer participants in the diverted analysis group thought that they had engaged in analysis or anticipated interest in identification of the animals than did the participants in the other two groups.

None of the individual analyses showed a significant difference between the memorize and the explicit induction groups. However, we cannot conclude that the memorize group was just as analytical as the explicit induction group. Aside from the usual difficulty of concluding that there was no difference, the performance of the memorize group was numerically closer to that of the diverted analysis group than was that of the explicit induction group on each of the measures. We expected that the memorize group would be similar to the explicit induction group, because both groups had nothing else to think about than the identification task. In retrospect, the memorization condition in Experiment 1 was not quite as indirect as we intended. Because the instructions were to “memorize which ones are blebs and which are ramuses,” the participants may have focused on category membership. If the instructions had, instead, told the participants that they would take a recognition test on the items and that they would have to be able to distinguish each item from a very similar one, the results might easily have been more implicit. Regardless, the present results indicate that a memorize task in the form we used has less of a claim to inducing nonanalytic behavior than does a diverted analysis task.

Overall, the hypothesis behind this experiment was supported. Despite excellent identification performance by all three groups, the diverted analysis group showed fewer signs of categorical analysis during acquisition and a greater tendency to believe that the category was simpler than it is than did either the memorize or the explicit induction group. In this, the diverted analysis procedure is a better simulation of beliefs about and performance with everyday categories than is either of the other procedures.

EXPERIMENT 2

Breeds Within Species

In Experiment 2, we tested the effect of diverted analysis training, using stimuli that had considerable perceptual variety within categories. The local style stimuli used

in Experiment 1 were deliberately developed to produce considerable perceptual uniformity within each category. However, the *breeds-within-species* (BWS) stimuli, shown in Figure 3, were modeled after categories, such as dogs and cats, in which there is considerable consistency within breeds but considerable variety between breeds within the same species. Using these stimuli, we asked our test questions about two *species*, represented by the two columns in Figure 3. To answer “yes, there is a feature in common to all blebs,” for example, the participants had to generalize across the perceptual variation represented by the different *breeds* shown in the top and bottom portions of each species column. The different breeds were presented to the participants as being long-range and short-range animals that differed in whether they could move without limit in a given direction or could move only two squares in that direction. That is, the long-range blebs could move any number of squares along a diagonal, but the short-range blebs could move only two squares along a diagonal in a single move. Using the analogy of dogs versus cats, this distinction in moves was intended to show that the behavior of the animals was similar, but not identical, within a category and very different across categories. The participants had to attend to the breeds to get the correct answer. However, they also had the analogy of moves between the short- and the long-range animals within the same category to support their seeing them as being in the same overall category. Consistent with the local style stimuli used in Experiment 1, these stimuli did not have any instantiations of features that appeared in more than one category and each instantiation within a category occurred on more than one item. Finding a simpler-than-it-is result with the diverted analysis procedure for the BWS stimuli was important because it would show that there can be perceptual variety within a category without eliciting analysis for consistencies within the category as a whole. Without this, the case for diverted analysis as an important condition for the learning of basic-level categories would be restricted to very homogeneous categories.

This experiment was a robustness check in another sense as well. In the BWS materials, shown in Figure 3, there was no equivalent to the low-salience dimensions that were present in the local style materials used in the previous experiment (e.g., front–back symmetrical/asymmetrical body shape or coarse/fine markings). That is, there was no feature that we saw or that was discovered by any participant that all blebs had and that no ramuses

Table 4
Responses to the Anticipation Question in Experiments 1 and 2

Experiment	No	?	Yes	% Yes
1A. Diverted analysis–local style	16	2	6	27
1B. Memory–local style	10	4	10	50
1C. Analyze–local style	8	1	15	65
2. Diverted analysis–breeds within species	21	1	2	9

Note—The question was, “Did you expect to be asked about the physical features of blebs and ramuses?” The “?” column records answers from the participants who said that they were not sure. The “% Yes” column records the percentages of “yes” responses among the participants who gave either a “yes” or a “no” response.

had. If the results of the previous experiment were to be replicated, they would not be dependent on the existence of such dimensions. The second purpose of the present experiment, then, was to show that the very special property of the local style stimuli—namely, the existence of perfectly predictive but low-salience dimensions—was not critical to the results shown in Experiment 1.

Method

Participants. The participants were 24 McMaster students, taking an introductory psychology course and participating for course credit. The participants were tested individually, and all were run in a diverted analysis condition.

Materials. The materials included the animals shown in Figure 3, with the chessboards, pairing, and sequencing being generated according to the criteria given in Experiment 1. Since there were two extra animals in this set, all the trials were increased accordingly. The informational descriptions of the animals in Figure 3 are given in the Appendix.

Procedure. The initial instructions were identical to those used for the diverted analysis group in Experiment 1. After these instructions had been given, the participant was given a response form on which to record the first initial of the species category and the number of moves of the winner (e.g., R 4; or the participant could write down the number of moves for each creature and then circle the least). As the participant completed each response, the experimenter quickly replaced the chessboard with the next. After Chessboard 3, the first short-range animals were shown. The experimenter explained that there were “short-range” bleeb, which could move only two squares at a time on the diagonal, and “short-range” ramuses, which could move only two squares at a time, either forward and back or sideways. The experimenter demonstrated movement patterns using Chessboard 4, leaving the bleeb and ramus in sight. The experimenter noted that the short-range bleeb and ramuses were smaller than the long-range ones, and that the participant could tell because they took up less room in the surrounding rectangle. The experimenter asked the participant to demonstrate the correct movement pattern of the short-range creatures on Chessboard 4. The experimenter corrected any movement pattern errors as they occurred and encouraged questions. If the participant made any inquiries about the category of the creature, the experimenter instructed, “If you’re not sure just guess and proceed with the moves.” The experimenter gave the participant feedback at the end of each trial: If correct, the experimenter said “Good,” removed the chessboard, and presented the next one. If incorrect, based on the identification of the creature, the experimenter said, e.g., “That’s a short-range bleeb and they move diagonally, 2 squares at a time. Please work out the chessboard again!” The experimenter also corrected moves at the end of each trial, including showing that there was a route that involved fewer moves. This was to encourage the participant’s sense that the chessboard task was the central task. For Trials 25–37 (one-at-a-time, diverted), the experimenter told the participant that he or she should complete the chessboard task as quickly as possible, still being accurate, and that he or she would be timed.

The test phase was identical to that used in Experiment 1.

Results and Discussion

Table 1 shows the percentage of errors on the final 12 trials, the trials on which single animals were presented. If, during these trials, the participants had been following a single-feature rule based on the high-salience dimensions, they would have made 16.7% errors. The error rate on these trials was 2.4%. No participant made more than one error.

Table 2 shows the average number and percentage of participants answering “yes” to the commonality ques-

tion, “Apart from the way they move, do you think that there is at least one feature or physical characteristic that all bleeb have exactly in common that no ramuses have?” For an answer to be recorded as *yes*, the participant also had to give a rule that named a single feature as sufficient, without an exception being mentioned. Because there were no perfectly consistent dimensions in these stimuli, no participant discovered a consistent dimension, so there was no equivalent here to the conservative and selective analyses in Experiment 1. The diverted analysis BWS group in this experiment answered “yes” to the commonality question 79% of the time. In comparison with Experiment 1, this was significantly more often than the explicit induction group ($\chi^2 = 10.243, p < .01$), but not significantly different from the diverted analysis group ($\chi^2 = 0.600, n.s.$).

The diverted analysis BWS group answered “yes” to the “did you analyze?” question 30% of the time (Table 3), significantly less often than both the explicit induction group ($\chi^2 = 20.299, p < .001$) and the diverted analysis group ($\chi^2 = 4.098, p < .05$) in Experiment 1. Finally, the diverted analysis BWS group answered “yes” to the anticipation question 9% of the time (Table 4), significantly less often than the explicit induction group ($\chi^2 = 15.769, p < .001$), but not the diverted analysis group ($\chi^2 = 2.655, n.s.$) in Experiment 1.

Thus, the diverted analysis BWS group replicated the pattern of results for the diverted analysis group in Experiment 1. Evidently, neither the overall perceptual homogeneity nor the existence of low-salience consistent dimensions is critical to the simpler-than-it-is effect. There is no way to compare the perceptual variety of these materials with that of natural categories, but at least some variety can exist without changing the effect of the diverted analysis procedure.

EXPERIMENT 3 Judgments of Commonality for Natural Concepts

In the introduction, we stated that most introductory psychology students are perfectly capable of identifying common objects but have oversimplified ideas about the criteria that allow them to do so. Although we are quite comfortable with the assumption that these students can reliably identify members of such categories as dogs, tables, and bottles, we thought that it would be best to demonstrate that they would answer “yes” to our commonality question. The aim of this experiment was to confirm that introductory psychology students will assert that common everyday categories have simple, perfectly discriminative features. The fact that introductory psychology students will answer affirmatively to such a question has been documented before (e.g., McNamara & Sternberg, 1983), but there were enough differences in procedure to make it advisable to demonstrate the result by using a question closer to the one that we used in the other experiments in this article. Specifically, we asked our participants, “Do you think that if you thought about it you could find *some feature* that would allow you to always tell that a

new item is a member of one of the following categories? That is, in the presence of a new object, could you always tell whether it is a member of the category by looking for this feature? For these categories, please give me your best guess about what such a feature might be.”

The categories about which we asked this question are shown in the leftmost column of Table 5. We used a wide variety of categories (artifacts, natural kinds, superordinates) in order to demonstrate that the belief in perfectly discriminative features is quite robust.

Method

Participants. The participants were 25 McMaster introductory psychology students, participating for course credit. They were tested individually at the end of other experiments.

Materials and Procedure. The participants were given a sheet that listed the six words shown in Table 5. They were then asked the question quoted in the introduction to this study. To corroborate our judgment that the features that they gave were insufficient to discriminate members from nonmembers of these categories, the same participants were then shown a sheet showing the items in the rightmost two columns of Table 5 and were asked, “Do you think that the feature that you named includes each of the items in the first part of the line and excludes the ones in the second part?”

The participants were asked to answer on another page, which contained the same six categories. Ample time was allowed for completion of both of the questions.

Results

The percentage of participants initially answering “yes” to the question about the existence of perfectly discriminative features, shown in Table 5, was dramatically higher than the proportion of participants who defended their choice of features in the face of the later include/exclude challenge. Only 6 of the 25 participants defended any of their criteria as being adequate to the challenge for any of the categories, and no participant defended his or her criterion for each of the categories for which he or she had said “yes” to the commonality question.

Finally, the percentage of people answering “yes” to the existence of perfectly discriminative features appears to be lower for the two categories that are clearly superordinate (fruit and furniture). However, the main result is that the ratings for all of these categories are substantially above zero: No participant denied the existence of single, perfectly discriminative features for all six categories. We

take this to corroborate the informal assertions made in the introduction about belief in the existence of simple, perfectly discriminative features.

GENERAL DISCUSSION

People show a persistence in seeking rules in standard laboratory tasks that, if shown in the world, should make them quite unlikely to believe that natural categories have simple defining features. Yet, as has been demonstrated in Experiment 3, among other places (and predicted in Medin & Ortony, 1989), people erroneously believe that such features exist for a wide variety of natural concepts. Clearly, something is wrong with the standard classification paradigm for generating this simpler-than-it-is belief; people in the lab seem entirely too analytic about the basis of categorization. The possibility investigated in this article is that analysis can result from the fact that most laboratory tasks give participants little else to think about other than the identification task itself, unlike most situations in the world. To remedy this, we put *identification in the service of use* by diverting the participants’ analytic activity to a task later in the action chain.

In Experiment 1, we demonstrated that the people in the diverted analysis group showed fewer signs of analyzing the basis of categorization than did the people in groups instructed to induce a categorization rule or memorize the stimuli. Specifically, they were more likely to mistakenly name features as perfect predictors, when such claims could easily have been eliminated if they had been evaluated in the presence of the stimuli. They were also less likely to report that they had been analyzing during acquisition or to have anticipated being asked about identifying the categories. In Experiment 2, we demonstrated that the effect also occurred with BWS stimuli, in which there was much greater perceptual variety among stimuli within each category than had been the case in Experiment 1. We would now like to discuss the conditions that we think are important for producing the simpler-than-it-is belief.

Important Acquisition Conditions

We conjecture that the less analytic category learning we observed was dependent on two aspects of our procedure. The first is *diverting analysis away from identification*: We

Table 5
Percentages of “Yes” Responses to the Commonality Question for Everyday Categories, Asked of the Participants in Experiment 3

Category	% Yes	Would the Feature You Named	
		Include	Exclude
Table	84	coffee table, pedestal table	counter, bed, stool
Bottle	88	baby bottle, pill bottle	jar, glass, carton
Dog	64	Chihuahua, greyhound	wolf, coyote, fox
Tree	84	young tree, pine, palm	bush, vine, bamboo
Furniture	56	dishwasher, wastebasket	table saw
Fruit	40	melon, strawberry, grape, coconut, pineapple	squash, pumpkin, almond

Note—The question was, “Do you think that if you thought about it you could find *some feature* that would allow you to always tell that a new item is a member of one of the following categories? That is, in the presence of a new object, you could always tell whether it is a member of the category by looking for this feature?” The last two columns contain the terms used in the challenge phase of the study.

focused the participants' attention on the consequences of making a categorization, rather than on an irrelevant task, as is done in most divided attention tasks. In this, the diverted analysis procedure is ecologically quite different from a divided attention procedure in which two tasks are seen as competing, rather than as different, aspects of the same overall task. However, focusing people's attention on subsequent use is not the only way of diverting their attention away from the identification task. For example, it also seems to be sufficient to induce learners to think of items as integrated wholes (e.g., by using orienting questions such as, "Would you like to have this animal as a pet?" or "Do you think this animal would thrive living in a swamp?"; Hannah, Skye, & Brooks, 2007).

The second aspect of our procedure that we think is crucial for producing less analytic category learning is *easy categorical identification during the acquisition phase*. In most circumstances in the world, people are free to attend to the identification task if it gives them problems. In the present experiments, one design factor that was included to avoid a perception of difficulty was having an initial period in which labels were provided, a learning support not uncommon in the world. However, we also believe that it is critically important that the stimuli provide many feature instantiations that are unique to a particular category—that is, have a *supportive stimulus structure*. In neither set of materials in this study did the same feature manifestations occur for more than one category (i.e., cue validity of 1.0 for every instantiated feature). Both sets of materials had feature instantiations that occurred in more than one item within a category (i.e., categorical validity of .8 for the instantiations in Experiment 1 and .33 or .5 for the instantiations of the breeds used in Experiment 2) and had multiple features with this highly discriminative structure. As was argued in the introduction, we think that many concrete categories in the world have all of these characteristics. Furthermore, we believe that such structures are critical for several everyday categorization phenomena.

In research with this kind of supportive materials, Brooks and Hannah (2006) showed that both people given explicit induction instructions and those given a rule in the form of a list of the relevant informational features for each category (e.g., "blebs typically have two legs, a rounded body, a rounded head and stripes") strongly relied on familiar instantiations in categorizing new stimuli. When these participants were given test exemplars that pitted multiple features that were rule consistent but had a novel appearance (e.g., a novel-appearing set of two legs, a novel-appearing rounded head, and a novel-appearing set of stripes) but also had a single familiar instantiated feature from another category (e.g., a familiar angular head), the majority of the participants classified the item on the basis of the single familiar instantiation. The participants thus allowed a single familiarly instantiated feature to override multiple informational features consistent with the rule they had been given or had induced. This reliance on specific instantiations of relevant features is a sensible policy in a world in which instantiations of relevant features are generally unique to everyday categories (e.g., the

legs of a baby or the legs of a hairy old man are uniquely human, a constraint also emphasized by Solomon & Barsalou, 2001, 2004), but the more abstract informational descriptors can often be applied to several categories (e.g., *two legs* applies to both humans and birds). As was argued in Brooks and Hannah, we believe that this tendency to learn and rely on specific instantiations is critical in understanding why so many everyday (and medical) rules are just feature lists that do not explicitly identify a decision procedure (e.g., "a bleeb has at least two of the following three features").

The experiments in this article held constant what we think are critical aspects of the supportive stimulus structure (no perceptual crossover of features between categories, same or similar instantiations of a feature on multiple items within a category, and multiple characteristic instantiations) and varied the acquisition task (diverted analysis, memorize, or explicit induction). An initial step in determining whether the various aspects of our supportive stimulus structure are, in fact, crucial is part of a paper in preparation (Jamieson & Brooks, 2007). In this research, a diverted analysis procedure was run on two sets of materials with the same informational structure as that used in the present Experiment 1, but differing in whether the same instantiated feature could occur in both categories. When the same instantiation occurred in both categories, the overwhelming majority of the participants reported analyzing during acquisition, but with category-specific instantiations, the diverted analysis results reported in the present article were replicated. However, there was no significant difference between the two sets of materials when the participants were instructed to analyze. Thus, for the analyze procedure, what seemed important was the informational description of the stimuli, whereas for the diverted analysis procedure, the surface manifestations were essential. Jamieson and Brooks also showed that it is necessary to have the same instantiated features occur in several members of the same category—that is, for the instantiations to have sufficient categorical validity.

Important Conditions for the Attribution of Perfect Predictors

We conjecture that for an attribution of coherence to occur, some level of support has to be present from each of two classes of variables. The first consists of *direct evidence from the task itself*. As one type of such evidence, the participants might check for discriminative consistency by performing an active, deliberate search of remembered items. They also might rely on a much less analytic impression of general consistency, which could be either memory for an impression gained during acquisition or a currently produced general impression. In addition, they could use a history of easy identification of novel items, implying a heuristic such as "the animals were all easy to identify; they must have something in common." In each case, the supportive structure of the present stimuli and, we conjecture, the stimulus structure of many everyday concrete categories give evidence of general consistency. The second class of variables consists of *background beliefs about the categories*. Since we used artificial materials,

our participants easily might have relied on laboratory-tailored rationales, such as “these animals have all been given a common move; this is a formal experiment . . . they must have something in common” (see Murphy, 2002, p. 129, for a discussion of this tendency). However, more generally, we might expect influences from any of a variety of general theories: Items that look alike rule alike; the physical attributes of a species are more likely to be consistent than its behavior; artifacts’ design characteristics tend to be consistent within a category; and so forth. Because of such beliefs, we might expect participants to be especially influenced by apparent consistency among features that are compatible with a plausible theory.

Our claim that the diverted analysis procedure is, in some measure, a simulation of everyday concept learning implies that we believe that the conditions just listed are true of many interesting everyday situations. Certainly, we do believe that people normally think about use or meaning, rather than identification, and that heavy contextual and structural support for identification is very common. People’s attitude about the obviousness of the identification of members of everyday categories suggests either that they have rarely experienced difficulty that they could not attribute to difficult viewing conditions or that any such experience has been swamped by many years of effortless categorization. In addition, there is ample evidence that people have background beliefs about natural categories that could either support or discourage attributions of coherence (see, e.g., the review in Murphy, 2002).

Backgrounding and Category Learning

The basic ideas underlying this work have been an important part of prior investigations. The effectiveness of directing people’s attention away from a specific task to another task in the same sequence of action was well illustrated in Jacoby, Levy, and Steinbach’s (1992) experiments on the effect of typeface on reading. When reading was made the focal processing task by having undergraduates reread a question, there was no effect of a shift in typeface (between an elite and a script face) from that used in the original presentation. When the reading was *backgrounded* by having people silently reread the question as part of answering the question, there was an effect of shifting the typeface on question-answering time. In both that work and the present experiments, making the task the focus of attention, rather than having it be a tool for accomplishing another objective, produced quite different results.

However, the present research is most obviously related to a very interesting task reported by Minda and Ross (2004), which they called *indirect learning*. In indirect learning, “learners are not told there are categories nor is learning about categories given as the goal (similar to unsupervised learning), but learning about the categories will lead to improved performance on the task on which they are being given feedback (similar to supervised learning)” (p. 1357). An example is an amateur photographer who might learn about different categories of light conditions and camera settings without ever receiving classification feedback, simply by seeing the good or bad

pictures that result. This indirect learning was simulated by asking learners to predict how much food they thought individual artificial animals should receive. Feedback on the correct amount of food was based on an unstated rule that depended on the size and category of the animal. The training items could be categorized on the basis of either a single, perfectly predictive feature or a family resemblance pattern of similarity established by four other features. Test items placed these two sources of identification in opposition, so that one category would be selected on the basis of the perfect feature and the other on the basis of family resemblance. Performance was compared between a group that merely made a prediction (the indirect-learning group) and a control group that first identified the category of the animal and then made a prediction. The prediction-only participants showed reliance on a greater number of features, even though a perfect criterial attribute was present. In contrast, the participants who first learned to classify the objects tended to learn the feature that was a perfect predictor.

Minda and Ross’s (2004) indirect-learning and our diverted analysis paradigms were developed independently of one another but are so obviously related that they merit close comparison. In both paradigms, learning about the categories is backgrounded to another goal to be accomplished with the stimuli. In both paradigms, a contrast is made between backgrounded learners and analytic learners who seem to set a priority on finding a single, perfectly predictive feature. Backgrounded learners fail to acquire information about the existence of perfectly predictive features that is acquired by explicit induction learners. In both paradigms, the structure of the items makes learning very easy, which is probably critical in letting backgrounded participants learn the categories without direct attention.

On the other hand, the two paradigms differ in the everyday situations to which they refer and which they are supposed to simulate. In indirect learning, the categories emerge due to their relevance to a focal decision and may not be explicitly named. This can lead to the learner’s showing surprise or to the reorganization of knowledge when the categories are focused on and explicitly named. In the referent situations for diverted analysis, the categories and names are well known and have been used for many purposes, and the surprise can come from the metacognitive discovery that these familiar categories do not have simple, common elements for identification or definition.

There are obvious ways in which these paradigms can reinforce one another. For example, it would be quite possible to make one dimension a perfect predictor in the diverted analysis materials, enabling test items similar to those in Minda and Ross (2004) and thereby allowing direct assessment of the contrasting information acquired by the diverted analysis and the analyze groups. On the other hand, it would be quite easy to add a metacognitive component to the indirect-learning paradigm. A general question such as, “Tell me about the animals. . . . What types were there?” could be used to assess whether the prediction-only learners would volunteer that there were two types of animal and whether they saw the materials as

consisting of animals that came in three sizes. Capturing the metacognitive aspects of learning is also an important part of understanding everyday conceptual knowledge, and these two paradigms clearly make distinctive contributions to this endeavor.

Simplified Personal Models of a Complex World

The main source of excitement for us in the present work is the suggestion that it might capture some part of the relation between learning to identify category members and developing a reflective model of the category—that is, capturing a way to provide enough space between the two processes to simulate the everyday discrepancy between them. The discrepancy between the complexity of the world and intuitions of simplicity has been observed in a wide variety of areas (to name a potent example, the discrepancy between intuitions that personality measures will be powerful predictors and the observed weakness of these measures; L. Ross & Nisbett, 1991). Certainly, the emphasis on uninterpreted features, dominant for many years in studies of category learning, was not conducive to studying this discrepancy. However, the current, very worthwhile emphasis on explanation-based learning (catalyzed by Murphy & Medin, 1985, and reviewed in Murphy, 2002) has led to designs in which surface features are tightly tied to the semantic interpretations being investigated. For example, in materials used by Murphy and Allopena (1994; and by Kaplan & Murphy, 1999, and Spalding & Murphy, 1996), every word or phrase given in the description of an item (e.g., *drives in jungles*) has a simple relation to a summary theme (e.g., *jungle vehicle*). This tight linking also does not promote the study of the discrepancies commonly observed between the features used for identification and the simplified, reflective opinions about the category. We believe that to simulate such discrepancies, it is essential to avoid the intense analysis that people bring to many standard laboratory tasks, even those tasks aimed at investigating explanation-based learning. What we hope the present paradigm has to offer is the possibility of investigating how people's simplified intuitions survive the necessity of acting in a complex world.

AUTHOR NOTE

This research was carried out with the support of the Natural Sciences and Engineering Research Council of Canada through a grant to L.R.B. Parts of this article were presented at the annual meetings of the Psychonomic Society, 1997 and 1998. We thank Lawrence Barsalou, Gregory Murphy, Andrew Neal, John Vokey, and Bruce Whittlesea for useful comments and discussion. We also thank Aaron Lautzenheiser for his contributions to an early version of this work. Correspondence concerning this article should be addressed to L. R. Brooks, Department of Psychology, McMaster University, Hamilton, ON, L8S 4K1 Canada (e-mail: brookslr@mcmaster.ca).

REFERENCES

- AHN, W.-K., & MEDIN, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, *16*, 81-121.
- BROOKS, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.
- BROOKS, L. R., & HANNAH, S. D. (2006). Instantiated features and the use of "rules." *Journal of Experimental Psychology: General*, *135*, 133-151.
- HANNAH, S. D., & BROOKS, L. R. (2006a). *Feature representation as a source of conceptual flexibility: Differential reliance on instantiated and informational features*. Unpublished manuscript.
- HANNAH, S. D., & BROOKS, L. R. (2006b). Producing biased diagnoses with unambiguous stimuli: The importance of feature instantiations. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *32*, 1416-1423.
- HANNAH, S. D., SKYE, A. L., & BROOKS, L. R. (2007). *Item integration and the failure of feature analysis in category learning*. Manuscript submitted for publication.
- HIGHAM, P. A. (1997). Dissociations of grammaticality and specific similarity effects in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *23*, 1029-1045.
- JACOBY, L. L., LEVY, B. A., & STEINBACH, K. (1992). Episodic transfer and automaticity: Integration of data-driven and conceptually-driven processing in rereading. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *18*, 15-24.
- JAMIESON, R. K., & BROOKS, L. R. (2007). *Supporting category learning: Enabling the illusion of categorical simplicity*. Manuscript submitted for publication.
- KALISH, C. W. (2002). Essentialist to some degree: Beliefs about the structure of natural kind categories. *Memory & Cognition*, *30*, 340-352.
- KAPLAN, A. S., & MURPHY, G. L. (1999). The acquisition of category structure in unsupervised learning. *Memory & Cognition*, *27*, 699-712.
- KEMLER-NELSON, D. G. (1989). The nature and occurrence of holistic processing. In B. E. Shepp & S. Ballesteros (Eds.), *Object perception: Structure and process* (pp. 357-386). Hillsdale, NJ: Erlbaum.
- MARKMAN, A. B., & ROSS, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592-613.
- MCNAMARA, T. P., & STERNBERG, R. J. (1983). Mental models of word meaning. *Journal of Verbal Learning & Verbal Behavior*, *22*, 449-474.
- MEDIN, D. L., & ORTONY, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-195). Cambridge: Cambridge University Press.
- MEDIN, D. L., & SCHAFFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.
- MEDIN, D. L., WATTENMAKER, W. D., & HAMPSON, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, *19*, 242-279.
- MINDA, J. P., & ROSS, B. H. (2004). Learning categories by making predictions: An investigation of indirect category learning. *Memory & Cognition*, *32*, 1355-1368.
- MURPHY, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- MURPHY, G. L., & ALLOPENA, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *20*, 904-919.
- MURPHY, G. L., & MEDIN, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289-316.
- PUTNAM, H. (1973). Meaning and reference. *Journal of Philosophy*, *70*, 699-711.
- REBER, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*, *118*, 219-235.
- REGHEH, G., & BROOKS, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *21*, 347-363.
- ROSCHE, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, 192-233.
- ROSS, B. H. (1996). Category representations and the effects of interacting with instances. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *22*, 1249-1265.
- ROSS, B. H. (1999). Post-classification category use: The effects of learning to use categories after learning to classify. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *25*, 743-757.
- ROSS, L., & NISBETT, R. E. (1991). *The person and the situation: Perspectives of social psychology*. Philadelphia: Temple University Press.
- SOLOMON, K. O., & BARSALOU, L. W. (2001). Representing properties locally. *Cognitive Psychology*, *43*, 129-169.

- SOLOMON, K. O., & BARSALOU, L. W. (2004). Perceptual simulation in property verification. *Memory & Cognition*, **32**, 244-259.
- SPALDING, T. L., & MURPHY, G. L. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 525-538.
- WARD, T. B., & SCOTT, J. (1987). Analytic and holistic modes of learning family-resemblance concepts. *Memory & Cognition*, **15**, 42-54.
- WHITTLESEA, B. W. A., BROOKS, L. R., & WESTCOTT, C. (1994). After the learning is over: Factors controlling the selective application of general and particular knowledge. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 259-274.
- WITTGENSTEIN, L. (1953). *Philosophical investigations*. New York: Macmillan.
- YAMAUCHI, T., & MARKMAN, A. B. (1998). Category learning by inference and classification. *Journal of Memory & Language*, **39**, 124-148.
- YAMAUCHI, T., & MARKMAN, A. B. (2000). Inference using categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 776-795.

APPENDIX

The Informational (Abstract Feature Description) Structure of the Stimulus Sets Used in the Experiments in This Study

(A) The informational structure of the salient dimensions of the local style stimulus set used in Experiment 1. The abbreviations are Bdy = body (0 is rounded, 1 is angular), Mrk = markings (0 is striped, 1 is dotted), Hd = head (0 is rounded, 1 is angular), Lgs = legs (0 is 2 legs, 1 is 4 legs).

Bleebbs				Ramuses			
Bdy	Mrk	Hd	Lgs	Bdy	Mrk	Hd	Lgs
0	0	0	0	1	1	1	1
0	0	0	1	1	1	1	0
0	0	1	0	1	1	0	1
0	1	0	0	1	0	1	1
1	0	0	0	0	1	1	1

(B) The informational structure of the BWS stimuli used in Experiment 2. The abbreviations are Bdy = body (0 is rounded, 1 is angular), Lgs = legs (0 is 2 legs, 1 is 6 legs), Hd = head (0 is rounded, 1 is elongated snout), Mrk = markings (0 is striped, 1 is dotted), Tl = tail (0 is short, 1 is long). Some items had alternative values, labeled X, to reduce the number of "cross-over" values.

Bleebbs					Ramuses				
Bdy	Lgs	Hd	Mrk	Tl	Bdy	Lgs	Hd	Mrk	Tl
0	0	0	0	0	1	1	1	1	1
0	0	0	0	1	0	1	1	1	1
0	0	0	X	0	1	X	1	1	1
0	0	X	0	0	1	1	1	1	1
0	X	0	0	0	1	1	1	X	1
1	0	0	0	0	1	1	1	1	X

(Manuscript received April 20, 2006;
revision accepted for publication June 28, 2006.)