Journal of Experimental Psychology: General
2006, Vol. 135, No. 2, 133–151

0096-3445/06/$12.00 DOI: 10.1037/0096-3445.135.2.133

Instantiated Features and the Use of "Rules"

Lee R. Brooks and Samuel D. Hannah
McMaster University

Classification "rules" in expert and everyday discourse are usually deficient by formal standards, lacking explicit decision procedures and precise terms. The authors argue that a central function of such weak rules is to focus on perceptual learning rather than to provide definitions. In 5 experiments, transfer following learning of family resemblance categories was influenced more by familiar-appearing features than by novel-appearing features equally acceptable under the rule. This occurred both when rules were induced and when rules were given at the beginning of instruction. To model this and other phenomena in categorization, features must be represented on 2 levels: informational and instantiated. These 2 feature levels are crucial to provide broad generalization while reflecting the known peculiarities of a complex world.

Keywords: rules, decision making, instances, categorization, concept learning

While discussing categorization in a cognition course, we asked a group of undergraduates the following question: "How do you know that an animal is a bird?" After checking that she had not by mistake stumbled into a philosophy course, a student answered, "A bird sings, flies, and has feathers." With a knowing smirk, we pointed out that an opera singer with a feather boa who flies on an airplane would qualify as a bird according to that rule. The student's reaction was not that this is a rare nonmember that has all of the characteristics in the statement. Rather, the example was taken as picky and a bit stupid. The example did not embody what she meant by "sings," by "flies," or by "has feathers." She meant, of course, birdlike singing, birdlike flying, and the self-generated feathers of a bird.

If this is what the undergraduate meant, then her statement falls short of being a formal rule in several critical ways. In addition to not specifying how the terms are to be weighted to make a decision, the terms are evidently not to be interpreted as broadly as the natural language terms allow and indeed are meant to be special to the category being characterized. Notice that the problem here is not just that people do not have proper weights on the terms (as in an inaccurate multiple regression). However weighted, a clear instance of flying does not help if it is the wrong kind of flying. But if the terms are not independent of the category (if "flying" is supposed to mean "birdlike flying" and "singing" is supposed to mean "birdlike singing" without further specification), then they seem circular and rather useless.

The idea that such rule statements are useless is clearly not the opinion either of the undergraduates or of academic physicians who regularly offer similar statements in medical instruction. As an example, there is a skin disease, lichen planus, for which a typical identification rule is "a pruritic [itchy], papular [a papule being a small, solid bump on the skin] eruption characterized by its violaceous [violet] color; polygonal shape; and, sometimes, fine scale. It is most commonly found on the flexor surfaces of the upper extremities, on the genitalia, and on the mucous membranes" (Chuang & Stitle, Section 2). As with the undergraduates' statements, no decision rule is specified (the decision rules in the psychiatric classifications of *Diagnostic and Statistical Manual of Mental Disorders* [4th ed.; American Psychiatric Association, 1994] are a rare exception in medicine). Further, if we offer a polygonal bump that is a perfect hexagon, no experienced physician would be comfortable with calling it lichen planus regardless of color, size, or location. Again, terms in the rule seemed to refer to lichen planuslike papules, not just any polygonal bump. Notice that for the lichen planus rule, the problem in communication does not stem from its being a careless or ill-considered production, as might be true for the undergraduates.

Possibly our diagnosticians—the undergraduates describing birds or the dermatologists working in their area of expertise—are simply giving a shorthand version that could easily be expanded if necessary. This is possible for some statements, but we would guess that the problem generally is more fundamental. Consider the task of specifying two legs in a way that would qualify a set of two legs as being human legs. First, note that one would have to include items as diverse as a baby's legs, a sumo wrestler's legs, a young woman's legs, a hairy old man's legs, and the legs of a starvation victim. One would also have to avoid including stork or ostrich legs. Similar to the dermatologist looking at a perfect hexagon, one would be reluctant to include a novel production such as a filmmaker's android legs, regardless of one's characterization. Mentioning that a human has two legs is useful for differentiating a person from a goat or a toaster, but it is hard to think up further specification that does not degenerate into a long disjunction of special cases. Even if such an expansion were successful, the resulting tome would no longer serve the purposes of efficient communication or medical instruction. The lack of

Lee R. Brooks and Samuel D. Hannah, Department of Psychology, McMaster University, Hamilton, Ontario, Canada.

We thank Rosemary Squire-Graydon for many valuable ideas as well as exceptional assistance in executing this research. The research was supported by a Natural Science and Engineering Research Council of Canada operating grant to Lee R. Brooks.

Correspondence concerning the article should be addressed to Lee R. Brooks, Department of Psychology, McMaster University, Hamilton, Ontario L8S 4K1, Canada. E-mail: brookslr@mcmaster.ca

further specification, at least in some cases, is not just a matter of carelessness or shorthand.

We are still left with the problem that if the terms as given are not independent of the category, if a bird is characterized in terms of birdlike flying, then they seem circular. One solution to avoid this problem is to leave the features broadly defined but increase the number of terms (e.g., add two legs, having a beak, living in trees, etc., to feathers, singing, and flying) and decide on the basis of some combination of them, such as the decision procedure embodied in a multiple regression. Any one such feature may overlap with other categories, but a sufficiently large number of them can reliably distinguish valid members from those of neighboring categories. However, as suggested by the bird–diva anecdote, the people giving such rules generally do not seem to believe that one has an insufficient number of terms. Rather, those people imply that one is simply misapplying the terms he or she does have. In some communications in medical instruction, this implication is not a gentle affair.

Despite these apparent insufficiencies in the "rule," if a person actually needed the characterization of bird, if he or she did not know what a bird was but otherwise understood the language, he or she would find the undergraduate's rule quite useful. When faced with an item identified as a bird, the person would, as directed by the rule, observe the flying, the singing, and the feathers rather than the color, breathing, or relation to a bread box. With exposure to a sufficient number of examples, the person would gain some knowledge about what constituted birdlike flying and birdlike singing. This presumably would be analogous to what medical students do for a good portion of their days. In the case of the medical students, however, the adequate use of the verbally stated rule would also be critical to their performance on tests, rounds, and eventually in (at least) the communicative aspects of practice. If the bird-free undergraduates or medical students were to meet a quite unfamiliar member of the category, as is likely for beginners, they might still have to rely on the general language extensions of the terms in the appropriate rule. Presumably, however, it would be with less confidence than if the features in the new object were very similar to previously encountered instantiations of those relevant features.

We argue in this article that the preceding examples illustrate the need for more than one concurrent level of feature specification for modeling and research in concept learning. There is a tendency in many discussions and modeling of categorization to use one level of feature specification for too many jobs. In some discussions and models, features are represented as a small number of elements of broad scope. The apparent goal is to account economically for broad generalization. That certainly was the goal in our bird and lichen planus examples; the whole category was to be covered with 3–5 terms. It was also the goal in early discussions of prototypes and current work aimed at coordinating with prior knowledge (Murphy & Allopenna, 1994; Rehder, 2003; Rehder & Ross, 2001). However many features are mentioned, they are typically on the level of "four legs" rather than "golden retriever four legs" or "golden retriever legs similar to Cleo's." The more general features support the desired points with less fuss than would the long list of more specific legs necessary to cover the domain of dog. Clearly, however, the much larger number of more specific features would have done at least as well for the specific goal of categorization, regardless of how badly the long list would have served the goal of communication. In other discussions, the features mentioned are much more specific, often strongly perceptual in nature (Barsalou, 1999; Brooks, Norman, & Allen, 1991; Schyns, Goldstone, & Thibaut, 1998). The apparent goal in such discussions is to account for expertise, with the subtext that the world is a very complicated and interactive place. The more detailed representations of features are also essential to discussions of perceptual symbol systems (Barsalou, 1999; Goldstone & Barsalou, 1998). As has been well documented in the literature, features at both of these levels have an important role to play. What we hope to document in this article is that both levels of features control aspects of generalization in the same transfer task and that a concurrent consideration of both levels is critical to understanding some of the peculiarities of communications and transfer using everyday rules.

Specifically, we would like to argue that when people give *everyday rules* (lists of features), they commonly *are naming objects of perceptual learning, not giving sufficient criteria for categorization nor implying decision rules*, such as "majority rules" or "best two out of three." When learning from a feature list, people use the terms of the rule to provide foci of attention for perceptual learning. When identifying new items, they commonly are looking for perceptually familiar instantiations, not just informational matches (e.g., lichen planus-type polygonal papules, not just polygonal bumps). To account for rule use, and many other categorization problems, people have to characterize features at two levels. The first level consists of *informational* (broad scope, relatively abstract) features, and the second consists of *instantiated* features (more specific manifestations of an informational feature for a particular category). Each informational feature in a category is a superordinate (possibly one of several) for the set of instantiations acceptable for that category. These principles apply as much to expert rules used in the formal instruction of identification as to the casual rules we elicited from undergraduates who have never considered them before.

These proposals have much in common with a number of themes in the recent categorization and concept learning literatures. Markman and Ross (2003) have emphasized a transfer appropriate processing analysis. The particular task used in learning determines what is learned about the material, and what is learned determines which transfer tasks elicit good performance (see also Whittlesea, Brooks, & Westcott, 1994). This framework recommends close task analyses of what is learned as well as the investigation of the interaction among different learning tasks; recommendations that we hope to apply here. Barsalou (1999) and Solomon and Barsalou (2001) have demonstrated the critical role for cognition of perceptually detailed representations, representations that are situated in the context and processing of particular actions. Our instantiated features constitute such a perceptual symbol system and confirm the importance of not relying on abstract, amodal symbols as in traditional accounts of cognitive processing. Schyns et al. (1998) have demonstrated the value of thinking of features as categories in themselves, categories that are formed in response to the conditions of learning rather than being selected from a preexisting list of fixed features. Our emphasis on the formation and selection of instantiated features, rather than on considering similarity solely on a whole-item level, is an application of their approach. The current proposals could also be seen as a special case of Sloman and Rips's (1998) argument for the necessity of combining “the flexibility of similarity-based inference and the compositionality and certainty associated with rule-based inference” (p. 87). What we intend to contribute to these themes is an argument for the need to consider features at two levels concurrently as well as what we hope is a useful task analysis of deliberate category learning in adults.

The experiments in this article are designed to apply the principles motivated by the informal analyses in this introduction to the learning of artificial concepts (i.e., to concepts using terms that initially are not practiced and sometimes not authoritative). Experiment 1 shows that a familiar instantiation of a feature has a greater effect on the classification of a new item than does an informationally equivalent but perceptually novel instantiation. At the same time, there is also very good generalization to markedly novel items that suggests that people are also using a more general form of knowledge. Experiment 2 provides the same kind of evidence for familiar instantiations when a weak rule is given to the participants at the beginning of training. This demonstration is important for our argument because weak rules (feature lists with no specified decision rule) are a common component of instruction in medicine, mineralogy, bird identification, and other areas of formal instruction. Experiment 3 demonstrates that a rule controls perceptual learning by directing attention to some features in an array rather than others. This will support our argument that this control of instantiation learning is a key function of everyday rules. Experiment 4 demonstrates that learned instantiations can be effective even when there is variation among the instantiations of a feature in the training items. In all of the other experiments, the instantiation of a feature was identical in all of the training items containing that feature. This experiment shows that the effect of a familiar instantiation occurs even when there is reasonable variety in the manifestations of a feature. Experiment 5 shows that a familiar instantiation has a greater effect on the interpretation of a neighboring ambiguous feature than does an informationally equivalent but novel instantiation. This demonstrates an effect that may be important in medical diagnosis: A clear, perceptually familiar feature may help to interpret less clear features, and this effect is at least partly a matter of perceptual form and not simply a matter of an informational state. Finally, Experiment 6 shows that more perceptual information is learned from the training items than simply a set of instantiations of features. Some effects of perceptual specificity that in the past have been attributed to learning whole instances (e.g., Brooks et al., 1991) may in fact be due to learning the instantiations of the categorically relevant features. Experiment 6 demonstrates that there is still an additional role for learning whole instances or some other form of relational information.

Collectively, we hope that these experiments demonstrate that perceptual familiarity can be a part of rule application as well as part of the less analytic influence of instances. Whole instances affect categorization by evoking a direct association with the category. Instantiated features, as characterized here, are associated with the much more analytic process of relying on some features as being conceptually and predictively critical to making a categorization. Again, the major background issue is the relation between the informational and instantiated characterizations of these relevant features. We turn to a discussion of the relation between these levels of feature characterization in the General Discussion section after we have the data in hand.

General Method

In all of these experiments, training deviates from standard practices in classification studies in several potentially important ways. Training in laboratory studies of categorization commonly presents unlabeled single items followed by trial-by-trial feedback until some learning criterion is achieved. In our experiments, items are initially presented as labeled pairs consisting of one member from each rival category. They are then represented as unlabeled pairs, again from both categories. In the final presentation before transfer to new items, single items are presented in a randomized order that is used across all participants. We believe our use of contrastive labeled pairs to be more similar to common everyday learning situations than is the single, unlabeled item induction procedure common in this area. Formal instruction in classrooms or on the job site often involves the labeling of exemplars and the provision of contrasting items and categories. Even informal learning is amply supported by observing labeling done by peers. Items in contrasting categories are readily available in the world and often intuitively sought out. Given that our concern in these experiments is with the nature of everyday rule knowledge and the bases of categorization of common objects, these experimental conditions have sufficient ecological validity to be interesting.

In addition, this procedure is conservative for our hypotheses in that using contrasting pairs could be expected to facilitate the extraction of structural knowledge (Gentner & Medina, 1998; Markman & Gentner, 1993), which could lead to participants placing less weight on any particular feature. By facilitating the extraction of structural knowledge, we are likely to be working against finding an effect of single familiar features on categorization. Finding an influence of single instantiated features under these conditions, as is done in all of our experiments, speaks to the robustness of instantiated feature knowledge.

For all experiments reported here, participants were randomly assigned to experimental conditions. All participants spoke English as their first language and were run in cohorts of varying sizes. Only participants scoring at least 80% correct on both classification rounds involving the singly presented training items (before and after test) were analyzed and reported. We found no reliable between-groups differences in training performance for any of our studies. Alpha was set to .05 for all studies.

Experiment 1: Initial Demonstration of the Power of Familiar Instantiations

Experiment 1 was designed to provide an initial demonstration that familiar instantiations of features have a greater effect on the classification of novel items than do informationally equivalent instantiations of the same named features. Participants first learned to categorize items into two family resemblance categories, shown in Figure 1, with the items in the left column called *bleebs* and those in the right column called *ramuses*. The top member in each column is prototypical for its category, and the items below them are the *one-aways*, each of which differ from the prototype on a different, single dimension. A critical feature of these items is that a given instantiation only appears in one category. For example, although items with four legs appear in both categories, the instantiation of four legs in the bleeb category is not the same as the instantiation of four legs that is common in the ramus category. This was meant to be an elementary simulation of a common everyday constraint, mentioned in the human–bird legs example given earlier: Although both birds and humans have two legs, the legs never had the same appearance in the two different categories.

After stating whatever rule they had induced during training, the participants were asked to categorize a set of transfer items, half of which are shown in Figure 2. This figure shows only the bleeb test

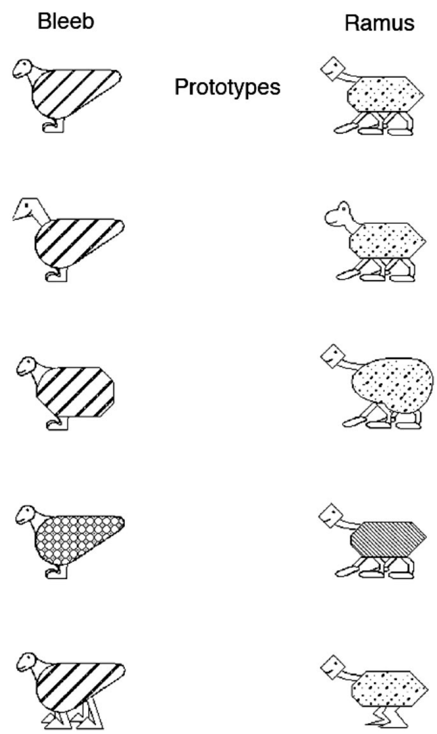

Figure 1. Training items used in Experiments 1, 2, and 5. The prototype of the category is at the top, with the one-away training items (deviant from prototype on a single feature) beneath each prototype.

items; during testing, these items were mixed with a comparable set of ramuses. At the top of Figure 2 are the prototypes of the bleebs and ramuses. In the left column are the nonprototypical bleeb training items, and in the middle and right columns are some far transfer items. These items correspond informationally with the training items shown in the left column but use different manifestations of the features. For example, the top items in the middle and right columns have a rounded head, striped markings, a rounded body, and four legs, exactly as does the training item in the top of the left (training item) column. These transfer items were designed to look quite different from the training items but be classifiable with the rules we used for the training stimuli.

The critical difference between the two sets (columns) of training items is that each of the items in the right column has one instantiated feature that is identical to an instantiation seen in the opposite category during training, whereas those in the center column have novel instantiations of the same informational features. For example, the top items in the center and right columns have four legs, atypical for bleebs, but the four legs in the right column are the four legs that had been seen on ramuses. Because all three of the novel features on these test items are informationally consistent with being a bleeb, this design pits the specific form of one feature against the informational value of three others. If, as hypothesized, the particular instantiation of the feature is important, the people categorizing the set in the center column should be more likely to call them bleebs than the people categorizing those in the right column. At the same time, because the items in the middle column have the same informational features as the training items and no conflict from instantiations associated with the opposite category, transfer should be good. That is, in these two groups, we should see evidence supporting the wide transfer we have attributed to informational features as well as the importance of the more specific transfer associated with particular instantiations.

Method

Participants

Twenty-eight undergraduate students at McMaster University participated for credit in a 1st-year psychology course. Half of the students were randomly assigned to the all-novel test condition, and half were assigned to the perceptual-interference test condition. These test conditions were run between participants to avoid the potential effect of the all-novel test items inducing a general informational feature strategy in the perceptual-interference condition.

Materials

Stimuli consisted of line drawings of imaginary animals displayed on a screen by a standard overhead projector. Participants indicated their responses on paper response sheets provided by the experimenter.

The training stimuli consisted of two categories of imaginary animals created around a family resemblance structure involving four features: head shape, torso shape, coat pattern, and number of legs. As can be seen from Figure 1, each category consisted of a prototype with all four features and four items that differed from their respective prototype by a single feature (one-away items). As can be also be seen in Figure 1, deviant features matched the characteristic feature of the other category only in terms of its informational content but had a different instantiation. In training, then, deviant features in training represented an informational overlap with the rival category but not a perceptual overlap.

Test stimuli consisted only of new versions of the one-away items. In the all-novel condition, all features were perceptually novel instantiations of the same informational structure used in training, as shown in the center column of Figure 2. As shown in the right column of Figure 2, the deviant feature for the perceptual-interference condition was perceptually identical to the corresponding characteristic value encountered in training. Thus, the bleeb one-away item with a deviant head shape not only has an angular head but also has the head characteristic of training ramuses.

Procedure

Training. Participants learned to categorize the training items over the course of three blocks. In the first block, labeled pairs were presented twice, once matched structurally (head-deviant bleeb, head-deviant ramus, etc.) and once with a pseudorandomly chosen partner, subject to constraint that it be from the other category and not be structurally identical. In this block, participants were not required to make a response. In the second training block, items appeared without labels, and participants were required to identify both members of each display. The experimenter provided the correct answers after all participants indicated they were finished. In the final block, single unlabeled items were shown in a random order. The experimenter again provided the correct answer when all participants indicated they were finished. For all blocks, the same ordering of stimuli was used across all participants in the experiment.

Test. Participants identified test items as bleebs or ramuses. The experimenter presented items individually in a single random order used across all participants in both conditions. Each item remained on display until all participants indicated they were done. After identifying the test items, participants indicated whether they believed that there was a single necessary and sufficient feature pertaining to bleebs and if so what it was. All participants were asked to provide a rule for classifying bleebs and

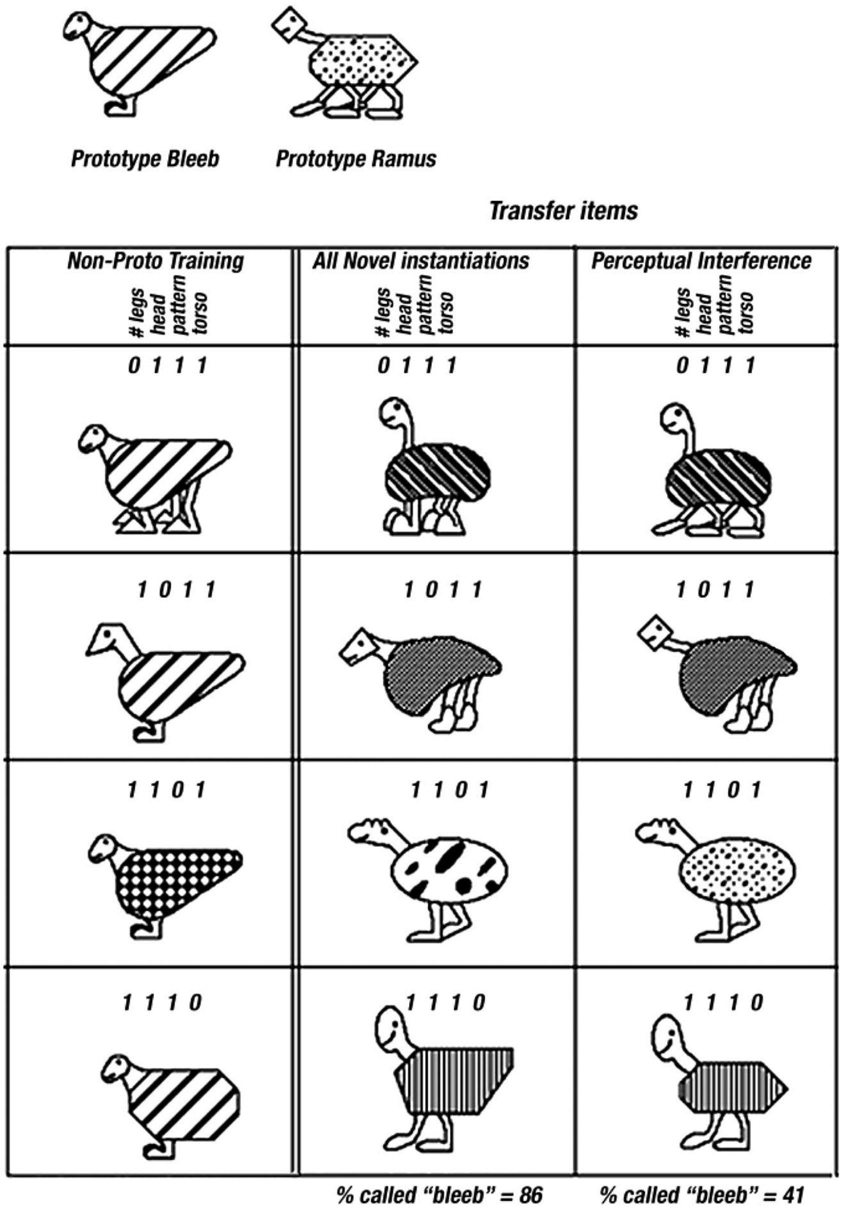

Figure 2. Training and transfer items used in Experiments 1 and 2. The left column, together with the prototype, shows the training items for the bleeb category. The middle column shows the four transfer items that should be called bleebs by the all-novel transfer group. This set of items has the same informational structure as the training items but has a novel instantiation of each feature. The items in each of the four main rows all have the same informational description. The right column shows four of the transfer items for the perceptual-interference group. For each of these items, the one nonprototypical feature had an instantiation that had been seen in the opposite category during training. The *1*s and *0*s associated with each item are the informational description of that item.

ramuses. Last, participants identified the training items, which the experimenter presented in a new random order, kept constant across all participants in the experiment.

Analysis

Mean differences in accuracy, scored according to the three-out-of-four classification rule consistent with the family resemblance structure, were analyzed by a *t* test.

Results and Discussion

Participants receiving all-novel test items had substantially, and significantly, greater accuracy in classifying test items than those receiving perceptual-interference items, $t(18) = -4.69$ (*df* corrected for heteroscedasticity), $SE = 0.32$, Hedges's $g = 1.77$. All-novel participants had an accuracy rate of 79% (6.36 items; $SD = 12.6\%$). Participants receiving perceptual-interference test items had an accuracy rate of 42% (3.36 items; $SD = 27.1\%$). This difference shows that familiar instantiations of features have a greater effect on the classification of novel items than do informationally equivalent instantiations of the same named features. Obviously, the informational features alone are not sufficient to predict the transfer performance.

However, it is also true that the all-novel group performed very well on the transfer items, evidently using a transfer policy that correlated well with the rule we used to build the items and to score the results. At the least, this high performance sets a limit on the specificity of what participants had learned and indicates the need for us to explain very good performance on transfer items whose appearances differ considerably from the training items.

The rules given by the participants generally contained 2–4 of the 4 features used in generating the training items. Only 5 participants mentioned only 1 feature, and only 2 participants specified a combining rule ("It was a bleeb if it had most of their features"). One other mentioned that combinations of features were important, and a 2nd participant provided a conjunctive rule for bleebs and a counting rule for ramuses. The report of the 19 other participants, then, was a list of 2 or 3 features with no combining rule. This finding provides empirical support for our anecdotal observations that people most commonly produce feature lists when generating rules. Examination of any medical text, bird identification guide, or many other instruction manuals for areas requiring extensive perceptual expertise reveals many rules in these domains similar to those given by our participants: 3–5 features with no decision criterion.

However, the participants' names for the features were not always identical to the terms that we had in mind when designing the transfer set. This raises the potential problem that when the participants described the features differently than we did, then the transfer items might not in fact have the same informational structure as the training items. For example, if a given participant described what we termed "a rounded body" as a "tear-drop shaped body," then the test items, which did not include any such tear-drop shaped bodies, did not provide a novel instantiation of this feature. If this were true, then we would not have met our intended condition of having the same informational structure in test as in training items. The one familiar instantiation might have had a large effect because according to the interpretation of the participant, there would be only a few informational features that were the same as they had been in training. To demonstrate that this potential mismatch is not necessary for obtaining an effect of the familiar instantiated features, in several subsequent experiments (Experiments 2, 3, and 6) we gave the participants a "rule," and therefore the intended names of the features at the beginning of training. In this way, the participants were given a description of the informational structure we built into both the training and test items. The current experiment provides evidence for the effect of instantiated features when participants attempted to induce the rule as well as some evidence for the types of rules that they expressed. The subsequent experiments are necessary to demonstrate an instantiated-features effect when we had communicated names for the features.

Experiment 2: Learning Instantiations After Getting an Explicit Rule

The purpose of this experiment was to show an instantiated-feature effect under conditions of explicit instruction. At the beginning of training, participants were told the four features that constitute the family resemblance structure of one of the categories without specifying an explicit decision criterion (i.e., no statement such as "at least three of the four features" was given). As suggested in the introduction, this type of weak rule is similar to those given in formal instruction in domains such as medicine, mineralogy, and bird identification. Finding an effect of familiar instantiations after explicit instruction is obviously an important step in supporting our interpretation of everyday rules.

Method

Participants

Sixty-four McMaster University undergraduate students took part for credit in a 1st-year psychology course. Half of these participants were randomly assigned to the all-novel condition, and half were assigned to the perceptual-interference condition. A minimum of 30 participants for each condition was set because we expected that the explicit feature list given at the beginning of training would reduce the size of any familiar-instantiation effect obtained.

Materials and Apparatus

The only difference from the previous experiment was that a feature list describing the four characteristic features defining the family resemblance structure of the training and test materials was provided in training.

Procedure

Procedures differed from the previous experiment by the provision in training of a feature list mentioned above. This list was displayed by an overhead projector at the start of training, and the experimenter read out the descriptions as well. The experimenter told participants that "Bleebs usually have . . . ," read the description of the four features displayed on the feature list, and then did the same for the ramus category. This feature list was kept visible during the display of the first round of labeled training items. A verbatim transcription of the experimenter's feature list instructions is given in the Appendix.

Analysis

Analysis was identical to that used in the previous experiment.

Results and Discussion

Participants receiving perceptual-interference test items again showed reliably lower accuracy than participants receiving all-novel test items, $t(62) = -2.39$,[1] $SE = 2.1\%$, Hedges's $g = 0.61$. The familiar-instantiation effect was reduced to a difference of 10

[1] As there was borderline evidence of heteroscedacity in the percentage correct data, we also ran a t test assuming unequal variance. This produced nearly identical, and significant, results. Because of both the concerns over heteroscedacity and the fact that many participants scored over 70% correct, leading to a possible dependence of the variance on the mean in the proportion data, we also ran a t test on the arcsine transformed proportions. This also produced nearly identical, and significant, results.

percentage points (all novel: 91%, 7.31 items, *SD* = 12.1%; perceptual interference: 81%, 6.5 items, *SD* = 20.8%). This experiment demonstrates that the instantiations of particular features have an effect even when participants have been given an explicit rule of the form often used in formal instruction.

As Kaplan and Murphy (2000) have pointed out, it is unlikely that all of the features involved in a categorization are related to some knowledge structure. Their point is certainly applicable to medical materials, which regularly have diagnostic features that cannot be derived (at least by the diagnostician) from what is known about the underlying disease process. In their experiments, they demonstrated a learning benefit for a set of material generated around a theme and showed that this benefit did not come at the cost of slower learning times for features not related to that theme by prior knowledge (i.e., background knowledge did not have the effect of limiting attention to or of handicapping the learning of features not initially related to the background knowledge). By analogy, we have found that giving people a rule did not have the effect of limiting learning of the form of features given in the rule. We might have expected that giving the informational features to the participants would concentrate the participants' attention on that level of features and thereby eliminate an effect of the instantiated features. Consistent with Kaplan and Murphy's data that question the sufficiency of an attention account of the effect of prior knowledge, knowing the value of an informational feature did not prevent our participants from learning and relying on specific appearances.

In the work that is most directly relevant to the current article, Yamauchi and Markman (2000) argued that perceptual variability and similarity have different effects on transfer depending on task. They found, for example, that perceptual variability has little effect on performance in an inference task but provided great difficulty in a categorization task. In their categorization task, the learner has the problem that no feature perfectly predicts the category, and the categorical identity of exemplars is withheld from the learner. Yamauchi and Markman argued that this uncertainty regarding class identity for the to-be-classified stimulus forces learners to compare the stimulus with all other encountered stimuli. The high variability in surface appearance combined with the overlap of features produces such complexity in these comparisons that the search for a set of diagnostic features is drastically impaired, and the discovery of structural relations in the domain is disrupted. This disruptive effect of surface variability is further documented in Markman and Maddox's (2003) follow up, which demonstrated that even variability along nondiagnostic dimensions impairs classification learning. The impairment due to surface variability (varied instantiations, in our terms), noted by Markman and colleagues (Markman & Maddox, 2003; Yamauchi & Markman, 2000), resulted from conditions in which the feature variance within categories approached that of the variance between categories. In the real world, this is likely to occur only when categories are very finely grained, such as when discriminating among different varieties of Monarch butterflies or when recognizing different individual faces.

The impairment discovered by Markman and colleagues (Markman & Maddox, 2003; Yamauchi & Markman, 2000), therefore, can be read as resulting from their participants' inappropriately recruiting instantiated features when the association between instantiation and category identity is abnormally low, such that the task set for the participants is best served by using informational representations. The variation in Yamauchi and Markman (2000) was intended to be a source of difficulty, and indeed it was. However, in the situations we attempted to simulate, the variation is strongly confounded with category and is therefore a resource. The family resemblance structure that led to a rule plus exceptions or disjunctive rules with the Yamauchi and Markman materials could lead the person to differentiate on the basis of instantiations with our materials. This additional resource led to quite rapid learning in our situation. We suspect that the difference in material could change this part of the predicted interaction between inference and categorization tasks. This stands as an interesting area for further exploration.

Experiment 3: Weak Rules Direct the Learning of Feature Instantiations

A central claim we made in the introduction is that a major role of everyday rules is to name objects of perceptual learning, rather than to give sufficient criteria for categorization. Experiment 3 was designed to document this role of a weak rule (feature list) in directing perceptual learning. The items in Figure 3 can be successfully categorized by either of two feature lists. An item is a bleeb if it has at least two of the features in either list:

> *Rule A*: Rounded, whistle-shaped torso; two plump, stationary feet–legs; short, wide fan-shaped crest.

> *Rule B*: Egg-shaped head on short neck; dark, widely shaped stripes; short, feathery tail.

A participant was given only one of the two lists. If the list directed attention to the features named in the list, then the participant should learn more about the instantiations of those features than the instantiations of the features in the alternate list, even though they are equally predictive of the category. Using the same types of transfer items as in Experiments 1 and 2, we predicted that the familiar instantiations named in the rule given to a particular learner ought to have a greater effect on categorizations than the familiar but unnamed instantiations. Again, despite equal exposure to equally predictive features, instantiations of the features named in the list ought to produce a bigger effect on categorization.

Method

Participants

Twenty-two undergraduate students at McMaster University participated in this experiment and received credit in 1st- or 2nd-year psychology courses as compensation. The 22 participants were randomly assigned to two different training groups, differing in which features (those in Rule A or those in Rule B) were named in the experimenter's instructions.

Materials

Training. Training materials differed from previous experiments in that there were six potentially relevant features, divided into two sets of three features. One set (Feature Set A) consisted of torso, feet, and crest, and the second set (Feature Set B) consisted of head, pattern, and tail. Exemplars of the bleeb training set with an informational description beneath each item are shown in Figure 3. A schematic description of the

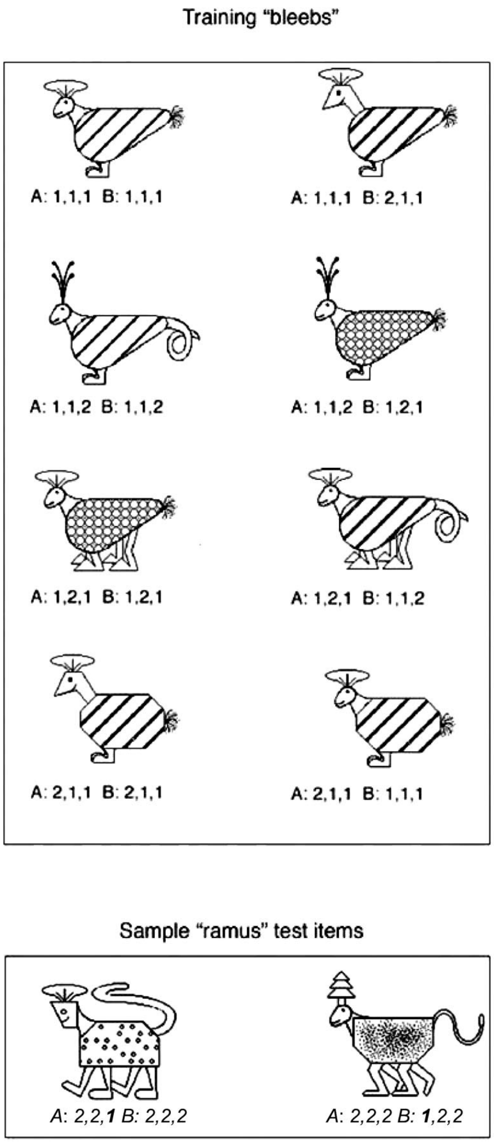

Figure 3. Top: Training items from the bleeb category used in Experiment 3. A comparable set of eight ramus items were mixed into this list for training. Features in Set A were named by the experimenter for approximately half of the participants, whereas those in Set B were named for the remaining participants. Beneath each item is its informational encoding. For Feature Set A: torso = round (1), angular (2); legs = two (1), four (2); crest = fan shaped (1), feathery (2). For Feature Set B: head = round (1), angular (2); pattern = stripes (1), dots (2); tail = short (1), long (2). Bottom: Examples of test items used in Experiment 3 for the perceptual-interference condition. Each feature is perceptually novel, except for one feature encountered previously among bleeb training items. The item on the left has a training bleeb's crest (Feature Set A), whereas the item on the right has the head (Feature Set B) seen previously on most training bleebs.

structure of the bleeb category is given in Table 1 (note that either Feature Set A or Feature Set B would equally well predict the category by a two-out-of-three rule). Although the nonprototype items were one-away items with regard to the set of features named in the rule given at instruction, most nonprototype items diverged from their respective prototypes by a total of two features because of the similar variation among the unnamed features (those in the unnamed set of features in the item). One exemplar in each category, however, deviated from the prototype only by a single feature. This was necessary to ensure that each feature took on a deviant value in two different exemplars.

Test. All test items deviated from the training prototype by a single feature. There were six test items from each category. For the perceptual-interference group, the deviant feature was a familiar instantiation originally seen in the other category. Perceptual interference came from a feature named in the rule given in training for three items in each category; perceptual interference came from an unnamed feature for the remaining three items. Examples from the ramus category are given in the bottom panel of Figure 3, depicting one item with perceptual overlap drawn from Feature Set A (left) of the training bleebs, and one with the perceptual overlap feature drawn from Feature Set B (right) of the training bleebs. It is worth noting that in this experiment, the perceptually interfering instantiation was pitted against at least two novel instantiations of named features and a total of five novel instantiations that were informationally predictive of the opposite category. For the all-novel participants, all features were perceptually novel, including the overlap features, regardless of whether they were named in the rule.

Procedure

Training and test procedures were identical to those in Experiment 1 with one exception. At the beginning of training, the experimenter pointed out three of the features from either Feature Set A or Feature Set B, noting that these would be useful in learning the categories. In addition to verbally mentioning them, the experimenter also listed them on an overhead. A verbatim transcription of the experimenter's feature list instructions is given in the Appendix.

Analysis

The data was analyzed by a 2 × 2 × 2 mixed-design analysis of variance (ANOVA), with naming (named, unnamed) as a within-subject factor and instantiation (familiar, novel) and counterbalance (Feature Set A, Feature

Table 1
Schematic of Category Structure for Bleebs

Training exemplar	Feature Set A			Feature Set B		
	Torso	Feet–legs	Crest	Head–neck	Pattern	Tail
Prototype	1	1	1	1	1	1
A	1	1	2	1	1	2
B	1	2	1	1	2	1
C	2	1	1	2	1	1
D	1	1	1	2	1	1
E	1	1	2	1	2	1
F	1	2	1	1	1	2
G	2	1	1	1	1	1

Note. Either Feature Set A or Feature Set B equally well predicted the category. Half of the participants received a rule naming Feature Set A, and half of the participants received a rule naming Feature Set B. Features coded as 1 are characteristic of bleebs, and those coded as 2 are characteristic of ramuses. As with all training materials in this article, the instantiations of features were never the same in the bleebs as in the ramuses. Torso: 1 = rounded and whistle shaped, 2 = six sided; Feet–legs: 1 = two plump and stationary, 2 = four thin and stepping; Crest: 1 = wide and fanlike, 2 = tall and treelike; Head–neck: 1 = egg-shaped head and wide neck, 2 = four-sided head and long, thin neck; Pattern: 1 = dark and widely spaced stripes, 2 = scattered dots; Tail: 1 = short and feathery, 2 = long and curved.

Set B) as between-subjects factors. Secondary analyses on the difference scores between named and unnamed features were also done.

Results and Discussion

Mean accuracy performance is summarized in Table 2. As in the previous experiments, these data show the interfering effect of a single feature that pointed to the opposite category than the remaining features in the item (novel > familiar). The unique contribution of this experiment is to show that this effect is stronger if the interfering feature is named in a rule provided to the participant. An interfering feature that was among those named in a particular rule produced lower accuracy than a feature that was equally predictive but not named in the rule (named accuracy = 80.2%, *SD* = 15.8%; unnamed accuracy = 97.2%, *SD* = 9.0%), $F(1, 38) = 67.09$, $MSE = 0.01$, $\omega^2 = .613$. A familiar instantiation of the interfering feature produced more interference than did a novel instantiation (perceptual-interference accuracy = 83.7%, *SD* = 17.4%; all novel accuracy = 94.2%, *SD* = 10.3%), $F(1, 38) = 14.06$, $MSE = 0.02$, $\omega^2 = .227$. Furthermore, the effect of naming was larger for perceptual-interference participants than for all-novel participants, with a significant Naming × Instantiation interaction, $F(1, 38) = 7.03$, $MSE = 0.01$, $\omega^2 = .217$. The three-way interaction of Naming × Instantiation × Counterbalance was marginally significant, $F(1, 38) = 3.68$, $MSE = 0.01$, $p = .062$, $\omega^2 = .182$.

Effects within each feature familiarity (perceptual-interference, all-novel) condition were analyzed with a pair of 2 × 2 mixed-design ANOVAs, both with naming (named, unnamed) as a within-subjects factor and counterbalance (Feature Set A, Feature Set B) as a between-subjects factor. For both analyses, the effect of naming was significant. For all-novel participants, this was the only significant effect, $F(1, 18) = 16.75$, $MSE = 0.01$, $\omega^2 = .437$, and they were less accurate on named features than on unnamed features (named accuracy = 88.5%, *SD* = 12.3%; unnamed accuracy = 100%, *SD* = 0.0%). For perceptual-interference participants, there was also an effect of naming, $F(1, 20) = 55.47$, $MSE = 0.01$, $\omega^2 = .714$. Again, the named features produced more interference than the unnamed, but the effect appears even larger (named accuracy = 72.7%, *SD* = 15.0%; unnamed accuracy = 94.7%, *SD* = 11.9%). There was also a Naming × Counterbalance interaction, $F(1, 20) = 5.86$, $MSE = 0.01$, $\omega^2 = .289$. This is not surprising in some ways. There is no reason why heads and feet should be equally salient. If heads were more salient than feet, they should be attended more in training than feet. If so, perceptually familiar heads would have a larger interfering effect at test than perceptually familiar feet. When unnamed, however, they go unattended, and no difference exists between unattended, or low attended, stimuli. Similarly, the greater learning of the more salient feature's appearance is of no importance when the appearances of all features are novel.

Because the order of effects of feature naming did not change with the interaction with counterbalance, we collapsed across this variable when examining the difference scores generated by subtracting the accuracy on unnamed features from that of named features for each participant. We also analyzed the full data set, with counterbalance included as a factor, and produced converging results. This analysis confirms that naming has a larger effect on the perceptual-interference participants than on the all-novel participants, $t(40) = 2.53$, $SE_{\text{pooled}} = 3.1\%$, Hedges's $g = 0.77$. Although the Hedges's g for the all-novel condition was much larger than for the perceptual overlap condition, it must be remembered that there was no variance at all for unnamed features in the all-novel condition, which drastically shrinks the pooled standard deviation for that effect. When we compared effects directly by looking at the difference scores, we saw that naming a feature reduced accuracy for perceptual-interference participants by 22 percentage points (*SD* = 15.7%) compared with neglecting features in instruction. Naming reduced accuracy for informational overlap participants, however, only by half that, 11 percentage points (*SD* = 12.4%).

The feature list evidently directed attention to the features it named, resulting in either greater learning of the instantiations or heavier weighting of them in test items. This result is consistent with Murphy and Allopenna's (1994) suggestion that meaningful features (by analogy for our case, the features named in a rule) should benefit from an attentional focus. However, the work of Spalding and Murphy (1999) and Kaplan and Murphy (2000) demonstrated that this prediction has to be made with a careful eye on the other coding strategies adopted by the participants.

By analogy to the birdlike flying examples given in the introduction, our participants seemed to rely on a "bleeblike, whistle-shaped torso" more than a novel instantiation that could also be called a whistle-shaped torso. The unique contribution of this experiment, however, was to show that the rule focused attention

Table 2
Mean Accuracy and Standard Deviations on Test Items in Experiment 3 (Effect of Feature Labeling)

Variable	Familiar features (perceptual interference)				Performance on familiar features		Novel features (all novel)				Performance on novel features		Overall performance	
	Set A		Set B				Set A		Set B					
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Named feature	.67	.14	.78	.15	.727	.150	.89	.13	.88	.12	.885	.123	.802	.158
Unnamed feature	.97	.07	.93	.15	.947	.119	1.0	.00	1.0	.00	1.0	.00	.972	.090
Performance across naming conditions	.817	.186	.854	.165			.945	.106	.939	.104				
Performance across naming and counterbalance conditions					.837	.174					.942	.103		

more on the instantiations of the features named in the rule than on other equally predictive instantiations. Of particular interest for supporting the perceptual learning aspect of this interpretation is that when no familiar instantiations were present at test, the effect of the rule in directing attention to particular features was drastically reduced.

Experiment 4: Similar But Not Identical Feature Instantiations

In all of the previous experiments, the same instantiations were used in all of the training items containing that feature in the same category. However similar the legs of one golden retriever are to the legs of another, it is surely not true that in general the characteristic instantiations of a feature in a category are literally identical to one another. In Experiment 4, the instantiations of the same feature in the training items were very similar but not identical to one another across all items. It is possible that when such variety of instantiations are used in training, the participants might rely on the informational features rather than being influenced by the one instantiation seen in training. The purpose of this experiment, then, was to show that physical identity of features across items within a category is not necessary to produce the effect of familiar instantiated features.

Method

Participants

Twenty-eight undergraduate students at McMaster University participated for credit in a 1st-year psychology course. Half of the students were randomly assigned to each counterbalanced instantiation condition.

Materials

Training materials were changed from those used in Experiments 1 and 2, such that the instantiations of the features were not identical across all of the training exemplars for the same category. Two sets of perceptual-interference test items were created using different instantiated features in training to counterbalance for feature salience. Bleeb exemplars used in training (left column) and in the perceptual-interference test conditions (middle and right columns) are shown in Figure 4.

Procedure

Training and test procedures were identical to Experiment 1.

Analysis

Data analysis was identical to that used in Experiment 1.

Results and Discussion

No significant differences in accuracy of classifying test items were found between the groups counterbalancing for the perceptual-interference feature used at test, $t(28) = -0.31$, $SE = 4.3\%$ (Counterbalance 1 = 55%, $SD = 18.2\%$; Counterbalance 2 = 54%, $SD = 15.1\%$). Therefore, these groups were collapsed into a single group for comparison with participants classifying all-novel test items.

A significant difference in accuracy of classification of test items was found among participants receiving perceptual-interference test items (55%; $SD = 16.6\%$) and participants receiving all-novel instantiations in their test items (80%; $SD = 3\%$), $t(58) = 6.15$, $SE = 2.0\%$, Hedges's $g = 1.59$. Therefore, interfering effects of instantiated features discussed in this article did not depend on there only having been one instantiation of a given feature in the training items.

Experiment 5: Biasing the Interpretation of Ambiguous Instantiations

This experiment was designed to determine whether a familiar instantiation has a greater effect on the interpretation of ambiguous features in an item, as well as of the item itself, than does an informationally equivalent but novel instantiation. If so, this would demonstrate an effect that may be of interest in many professional categorization tasks, including medical diagnosis: A clear, perceptually familiar feature may help to interpret less clear features, and this effect would be at least partly a matter of perceptual instantiation and not simply a matter of an informational value.

After the same training as in Experiment 1, participants were shown items such as those in Figure 5. Three of the features in each of these animals have a mixture of the contrasting informational features from training. For example, the animal in the top left of Figure 5A could be seen to have markings that are stripes but also as composed of dots. The legs could be seen as two but also as four legs in which two are nearly occluded. The body shape has both rounded and angular components. The fourth feature in each item in Figure 5A is the biasing instantiation, coming from either the bleeb or the ramus prototype. For example, the head on the top left item is the prototype bleeb head instantiation, and the head on the item to its right is the prototype ramus instantiation. If this clear feature biases the interpretation of the ambiguous features, then there should be more bleeb interpretations with the bleeb head and more ramus interpretations with the ramus head. We expected that the animal in the top left would be more likely to be interpreted as having stripes, two legs, and a rounded body, whereas the animal to its right would be more likely to be said to have dots, four legs, and an angular body. Figure 5B is the control for the effect of informational value. The heads in the top row are rounded on the left, and they are angular on the right, just as for Figure 5A, but in this case, they are novel instantiations. If the familiar instantiations are important, there will be more category-consistent identifications in Figure 5A than in Figure 5B. A given participant will be asked to identify the items in Figure 5A or the items in Figure 5B and to justify his or her decision by listing supporting features.

The measures were how many category-appropriate shifts of categorization a particular participant gives for each of the two panels and what proportion of participants reversed their interpretation of the ambiguous features listed in support of their decisions. Listing the same feature for both members of an ambiguous pair while reversing the classification of the second pair would indicate a reversal in the interpretation of that feature. For example, suppose a participant classified the top left test item in Figure 5A feature as a bleeb and indicated the ambiguous legs as a supporting feature. If the participant later called its matched item, the item in the top right of Figure 5A, a ramus but still put down legs as a supporting feature, this would indicate that the change of the

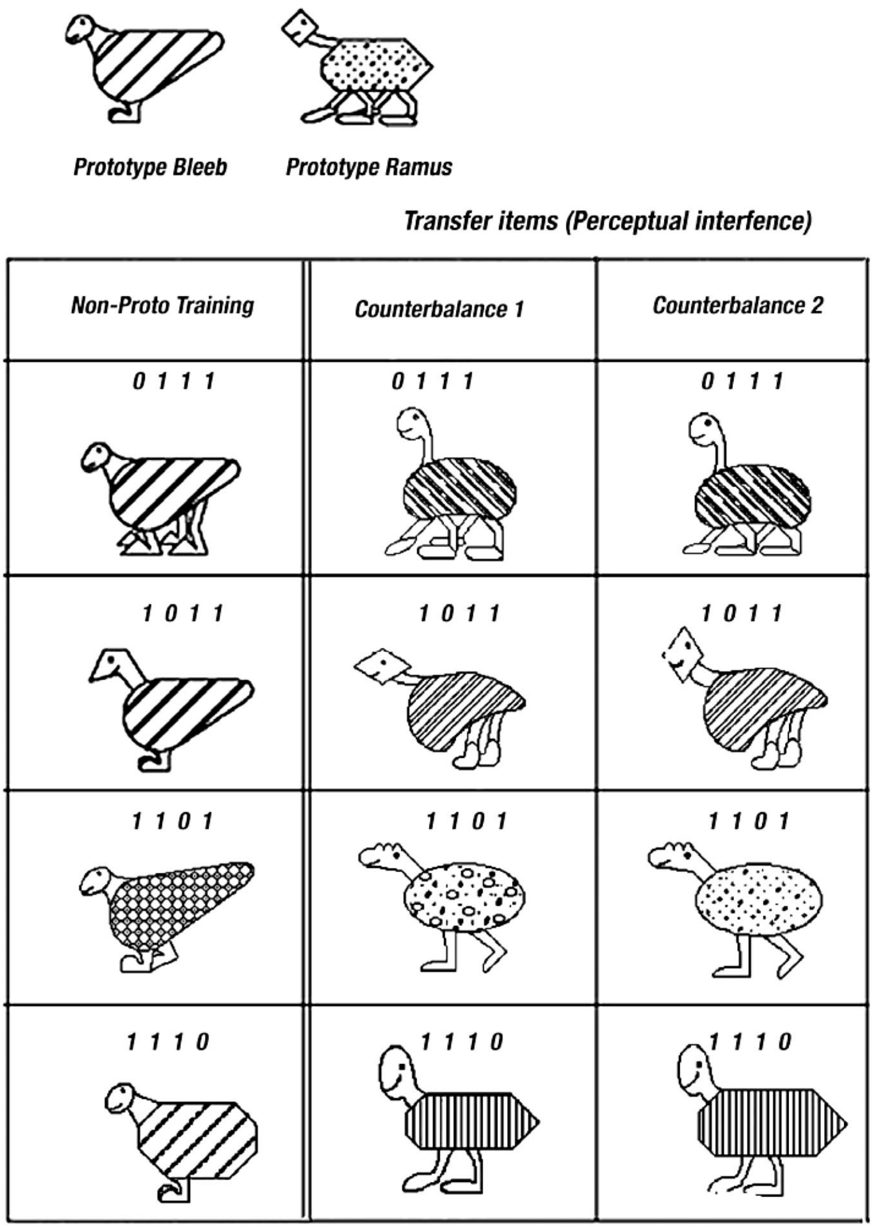

Figure 4. Training bleebs (left column) and test bleebs (middle and right columns) in the perceptual-interference conditions in Experiment 4. Each of the instantiations of the training features is a different variant of the prototype features.

unambiguous features (heads) changed the interpretation of the ambiguous legs.

Method

Participants

A total of 56 undergraduate students from McMaster University took part in the study, with 28 participants being randomly assigned to each transfer condition. Participants received course credit for either a 1st- or 2nd-year psychology course.

Materials

Training. Training items were identical to those used in Experiment 1.

Test. Test materials consisted of paired versions of the same four ambiguous test animals. For each of these items, two or three features contained a mixture of properties, such as heads that had both rounded and angular elements. In the familiar-instantiation condition, one version of each item had a bleeb training feature, and one version had a ramus training feature (see Figure 5A). In the novel instantiation condition (see Figure 5B), one version had a novel instantiation of a characteristic bleeb feature (e.g., novel-looking rounded head), and one version had the ramus equiv-

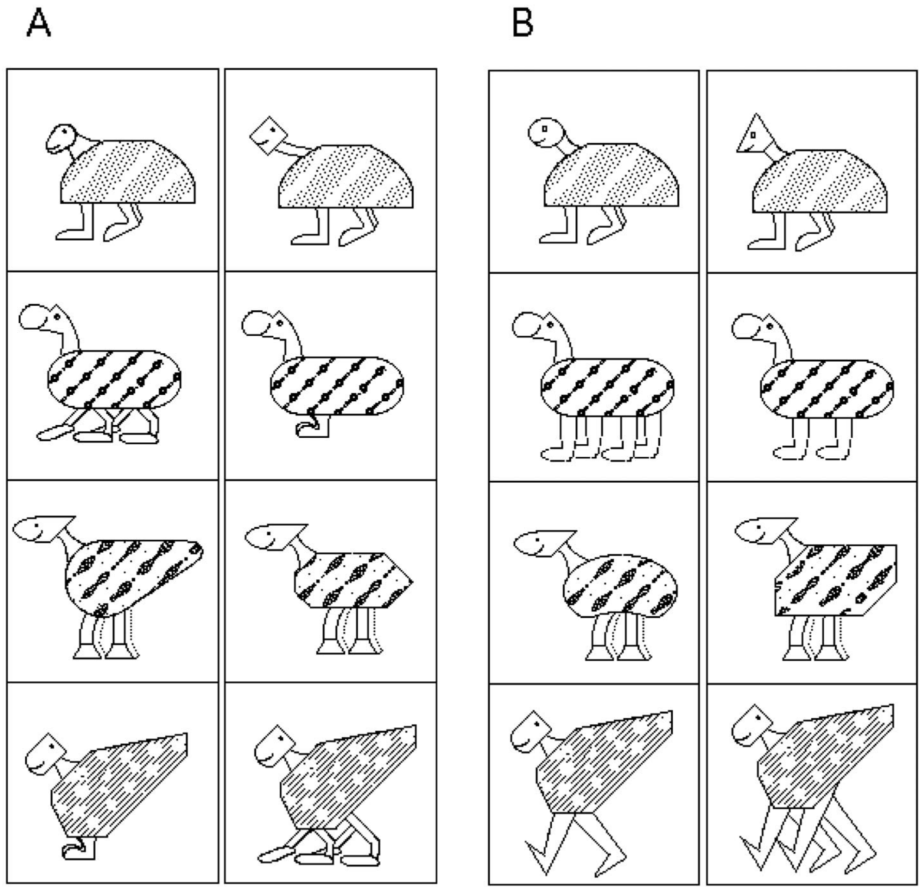

Figure 5. Familiar-instantiation test pairs (A) and novel-instantiation feature test pairs (B) in Experiment 5.

alent. A pilot study suggested that some participants found the ramus pattern somewhat ambiguous itself; thus, a second set of items with unambiguous feet were used instead of a pair of pattern-based items. In debriefing participants from the previous studies, and examining the features listed in rule statements, we discovered that although head and torso were highly salient, both feet and pattern were less salient and roughly equivalent in their saliency.

Procedure

Training was identical to Experiment 1. At test, either the novel or familiar feature instantiation materials were shown, one at a time. Participants saw both pair members, with items shown in a pseudorandom order, subject to the constraint that at least two items separate pair mates. As in the previous experiments, participants merely identified each of the eight items they saw. Which item was presented first was counterbalanced across participants. After all of the items had been identified, the items were individually represented in the original order, and participants were asked to justify their decisions for each item by indicating supporting features.

Analysis

The dependent variables in this experiment were the number of reversals of item classifications and of feature descriptions across paired stimuli. The difference between transfer groups was analyzed using a *t* test. A Mann–Whitney *U* test was also performed on the same data and produced convergent results to the *t* test reported here.

Results and Discussion

Participants receiving ambiguous test items with a familiar instantiation of the biasing feature were significantly more likely to reverse their classification as the biasing feature changed than were participants receiving test items with a novel instantiation of the biasing feature, $t(54) = 5.96$, $SE = 0.12$, Hedges's $g = 1.62$ (familiar instantiation = 2.21 reversals, $SD = 0.88$; novel instantiation = 0.82 reversals, $SD = 0.92$). When the biasing feature was perceptually novel, 23.5% of participants reversed their description of one or more ambiguous features to be consistent with the unambiguous feature's category. When the biasing feature was perceptually familiar, however, 42.9% of participants reversed their description of at least one ambiguous feature to be consistent with the unambiguous feature's category (percentage difference = 19.4). A test of the equality between conditions of the proportion of reversals confirms that they are reliably different ($z = -506.65$, normal approximation to the binomial). Evidently, the instantiations, not just the informational features, were influential in determining the classification of items with ambiguous features and the features themselves. We take this to be analogous to the phenomenon of a biasing piece of information changing the identification of items and features in medicine (e.g., LeBlanc, Norman, & Brooks, 2001). The difference is that in the current study, the differing interpretation of the item and the features occurred within the same participant within 1 min or 2. Further, the items were relatively simple and did not depend on a complex educational history.

Experiment 6: "Breeds Within Species"—Specificity Effects Beyond Lists of Instantiations

Experiment 6 was designed to show that more perceptual information is learned from the training items than simply from a set of feature instantiations. Some effects of perceptual specificity that in the past have been attributed to learning whole instances may in fact be due to learning the instantiations of categorically relevant features. For example, the diagnostic advantage of previously having seen a perceptually similar dermatological case might have been due to the similarity of isolated feature instantiations rather than to the overall similarity of the whole case (configural properties, correlated features—whatever makes two items look similar overall rather than just similar in one or two parts), as suggested by Brooks et al. (1991). If so, then the diagnostic rules might be exerting a more direct effect of perceptual specificity than we had supposed. This experiment was designed to determine whether there is still an additional role for learning whole instances or some other form of relational information.

The materials in Experiment 6 were generated by analogy to breeds that form clusters within species. In general, dogs differ from cats on a number of characteristics. However, within these species differences, there are clusters formed by different breeds, such as spaniels and terriers, Siamese cats and Persian cats. Despite the consistency of cats relative to dogs in shape of eyes, configuration of whiskers, movement of tails, and posture when sitting, the particular breeds have distinctive manifestations of these characteristics. Consequently, a picture of a cat with a Persian face on a Siamese body would look decidedly odd, possibly leading to the suspicion that the picture had been altered. However, that suspicion would not be based on anything that was incorrect about any single informational or instantiated feature. Instead, the rejection of the picture would be based on the combination of individually correct features. The materials in Figure 6 were designed to provoke this kind of judgment.

In the top panel of Figure 6, it can be seen that each of the two categories form a family resemblance "species" based on five informational features: roundedness of head, roundedness of body, number of legs, markings, and length of tail. However, the top three and the bottom three animals in each species form clusters of distinctive instantiations of those features (i.e., the instantiation of the same informational feature is identical for members of a single "breed" but differs from the instantiation of that informational feature for the other breed). The bottom panel shows a sample display in which a participant who had learned the training items might judge one of the items as consisting of an odd combination of individually familiar features—a nonbiological hybrid between Persians and Siamese. If this were indeed the judgment, then clearly the participants would have learned more about the items than the individual-feature instantiations. This knowledge of common association among features is one of the arguments that has been advanced in favor of instance models (e.g., Wattenmaker, 1993) and is not accounted for by the instantiated-feature hypothesis argued for in this article.

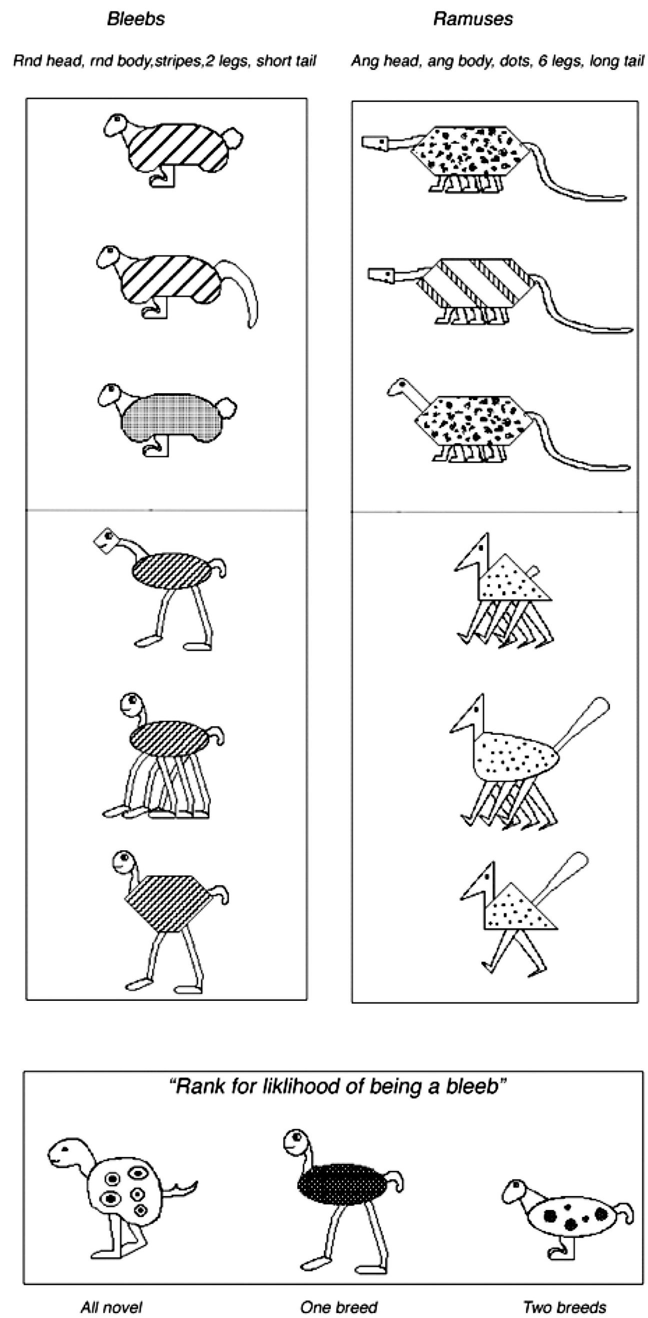

Figure 6. Top: Training items used in Experiment 6. The prototype for each category is depicted at the top of each column. The first three items in each column represent one "breed" within each "species," and the bottom three items represent a second breed within the same species. Bottom: An example of a choice presented in a test trial. All three items instantiate the same information embodied in one of the training one-away items. For the item on the left, all features are perceptually novel. For the middle item, the four category-consistent features (head, torso, tail, feet–legs) all come from the same training breed (bottom three items in the left column of the top panel). For the item on the right, the head and feet were seen in training in one breed (represented by the top three training items), whereas the torso and tail were seen in training in the other breed members. The labels *all novel* and so forth were not shown to participants.

Method

Participants

Sixteen undergraduates from McMaster University took part in this experiment for credit in 1st- and 2nd-year psychology courses.

Materials and Apparatus

Training. One characteristic feature, tail, was added to the four-feature structure used in most of the prior experiments. This resulted in each training category having six members, five one-away items and one prototype (see Figure 5). Most important, within each category, two subordinate categories were created on the basis of the feature values irrelevant to the basic-level distinction. Each species, therefore, consisted of two breeds. For example, although bleebs usually had rounded heads, rounded torsos, two legs, stripes, and short tails, for half of the items (first three in Figure 5), the torso was kidney shaped, the head was egg shaped, the legs were short, the tail was rabbitlike (when short), and the stripes were thick (when present). For the remaining items, the torso was usually oblong, the head was circular, the legs were long, the tail was hook shaped, and the stripes were thin and wavy.

Test. Test items were presented in triads. Each member of a triad was identical at a purely informational level, embodying the same information as the training one-away items. One member was composed of perceptually novel instantiations of a given one-away item's characteristic informational values (all-novel items). For the remaining two items, all of the category-consistent instantiated features were seen in training, and the one overlap feature was a novel instantiation of the overlap feature seen in training. Critically, these two items differed in that for one triad member, all four category-consistent features came from the same breed (one-breed items), whereas for the remaining item, two came from one breed and the other two came from a different breed (two-breed items). An example of a test triad is shown in Figure 6.

Procedure

Training was identical to that used in Experiment 1, with one important exception. Prior to displaying the training items, the experimenter read a short description of the characteristic features of the two categories: "Bleebs usually have a rounded head, a rounded body, two feet, a short tail, and stripes. Ramuses usually have an angular head, an angular body, six feet, a long tail, and dots." At test, participants ranked triads of items as to which was the best example of a suggested category. The suggested category was always the category consistent with the majority of features. Each item was ranked on a three-point scale, with 3 indicating the item least likely to be a member of the suggested category, and 1 indicating the item in the trio most likely to be a member of the suggested category.

Analysis

Mean rankings of the three types of test items were computed for each participant: one-breed items, two-breed items, and all-novel items. Differences in the mean rankings were analyzed with the nonparametric equivalent of the one-way repeated measures ANOVA (Friedman test), as the rankings are nonindependent. We also ran planned comparison on the difference between within-breed and between-breeds items using the nonparametric analogue of the paired t test (Wilcoxon signed-ranks test). We expected that one-breed items would be ranked lower (more likely to be a member of a given category) than two-breed items and that these in turn would be ranked lower than all-novel items. As it was not clear what was a meaningful measure of effect size for such nonindependent ranks, we present only the raw difference between mean ranks.

Results and Discussion

There was a significant difference in typicality ratings across test items, as revealed by a Friedman rank test of k correlated samples, $\chi^2_{F(2,N=16)} = 26.8$. The mean difference between the lowest ranked items (one breed) and the highest ranked items (all novel) was 1.2 rank steps (1.5 vs. 2.7). One-breed items were reliably considered more likely to be members of their respective categories than two-breed items (1.5 vs. 1.8), $T(12) = 13.5$, which in turn were reliably considered more likely to be members of their respective categories than all-novel items (1.8 vs. 2.7), $T(16) = 0$, *SE*s $= 0.06$ for all three groups. The lower ranking of the all-novel condition than the others shows that familiar instantiations contribute to judgments of category membership above that contributed by informational values, the main point of all of the other experiments in this article. The higher score for the one-breed condition over the two-breed condition shows that learners responded to the familiarity of instantiated-feature combination and not just the familiarity of each of the feature instantiations themselves.

However, prior studies have found little evidence for a role for correlated features. For example, Malt and Smith (1983) found little to no contribution to typicality judgments of naturalistic categories when given items consisting of correlated features over items with uncorrelated features once the cue validity of the features was taken into account. Similarly, Murphy and Wisniewski (1989) found no evidence for the learning of a perfectly correlated feature in an experimental study using experimenter-created materials. Chin-Parker and Ross (2004) found evidence that perfectly correlated features were learned when learning took the form of an inference task rather than an induction task, but even here, the effect was small.

The current study was worth running even in the face of these weak prior results because of our interest in the distinction between informational and instantiated features. The studies above use only feature structures that we would refer to as informational. We have argued throughout this article that the different levels of features play different roles in categorization. This distinction makes the current demonstration for instantiated features important for our argument.

A related—but different—distinction, that of verbal versus perceptually mediated features, might also explain why we have been able to show some learning of correlated features even though the features were more weakly correlated than many previous studies that have had difficulty finding such evidence (e.g., Chin-Parker & Ross, 2004; Murphy & Wisniewski, 1989). The features of verbal stimuli are highly discrete and parsed (constructed) for the participant by the experimenter. In contrast, even with simple, perceptually mediated materials in which the learning is accompanied by instruction regarding diagnostic features, the particular boundaries have to be determined by the participant. This necessitates attending to at least adjacent features concurrently, allowing incidental learning of feature co-occurrence to a greater degree than in studies showing little learning.

It is still true, however, that the effect of correlated features (relational or instance information) was small in this study. We conjecture that this is true because both the conditions and materials in all of these studies encouraged attending to separate features. This will be discussed more fully in the General Discussion.

For the present, we restrict ourselves to the conclusion that we cannot assume that previously reported instance effects (e.g., Brooks et al., 1991) are solely due to the effects of individual instantiated features.

General Discussion

In this article, we have argued that familiar instantiations of features play a vital role in categorization distinct from the informational features directly captured in rules and informal communications. In all of the preceding experiments, we have shown that a familiar instantiation has a greater, often markedly greater, influence on categorization than does a novel feature with the same informational value. This effect of familiar instantiations occurred both when the participants were attempting to induce the bases of categorization (Experiment 1) and when they were actually given a rule (Experiments 2, 3, and 6). Experiment 3 showed that this everyday type of weak rule (a feature list with no decision procedure) had the effect of directing attention to some features and establishing the instantiations of those features as especially influential in item identification. These familiar instantiations can also influence the interpretation of ambiguous features more than do novel instantiations with the same informational value (Experiment 5), and their effect was not restricted to the special case of there being only one instantiation of a feature within a category (Experiment 4).

In evaluating this evidence, we must first point out that we believe the normal role of a familiar instantiation is to augment and help interpret other evidence, not to reverse a decision based on all of the other features. Putting a familiar instantiation in opposition to several novel features is an experimental convenience, not a suggestion that a familiar instantiation normally contradicts less familiar evidence. In this light, the more realistic test condition would have been to put a familiar bleeb feature in the company of novel bleeb features and evaluate performance against a bleeb with entirely novel features. The problem, of course, is that, for these materials, performance with the novel features was close enough to ceiling to leave a very restricted range within which to observe our effects. A second special constraint on these results is that the test items were all sort of like new. The all-novel conditions provided items that were perceptually unfamiliar but that provided material to which the induced or presented rule could readily be applied. This is commonly the situation of the beginner, more rarely of the expert, faced with a strange presentation. Obviously, we would pick different transfer items if we were attempting to investigate the perceptual familiarity of whole instances. These items are very informative, however, for investigating the specialization that occurs in developing and learning to apply a rule.

The Task-Specific Value of Instantiated Features

Regardless of experimental tactics, these results raise the question as to why anyone should rely so strongly on a single perceptual match, in some cases more strongly than several novel instantiations of features that predict a different category. We argue that the major reason is that a single, perceptually rich feature is extremely diagnostic of the category. The informational feature of two legs leaves doubt as to whether the possessor is a human or a bird, but an instantiation as a baby’s legs removes question. An instantiated feature also can provide sufficient information under restricted or distorted viewing conditions. The head of a dog emerging from behind a fence leaves little doubt as to what will soon be visible; a familiar polygonal papule might allow diagnosis when the normal forms of other cues have been destroyed by scratching or rendered unfamiliar by being on the skin of a person of an unfamiliar race or age. Another way to state this constraint is to consider the within- and across-category distribution of such features. For many concrete, named categories, strong perceptual similarity of parts within categories is common. The legs of one goat are likely to strongly resemble the legs of at least some other goats. By contrast, strong perceptual similarity of parts across categories is comparatively rare—although informational identity of individual parts can be frequent. The goat legs are unlikely to strongly resemble the legs of any cow, despite both animals sharing the informational property of four legs.

The confounding of feature instantiation and category in our material also is critical to other phenomena. Hannah and Brooks (in press) have shown that categorical biasing effects—similar to those found in the medical cognition literature (LeBlanc et al., 2001)—critically depend on instantiations being highly predictive of category. In a companion article (Hannah & Brooks, 2006), we showed that effects of familiar instantiations, such as those shown in Hannah and Brooks (in press) and in the current article, are due to differential weighting of the familiar feature rather than to attention being distracted from novel features.

Considering the development of category-specific features, as documented by Schyns et al. (1998), our speaking of category-specific instantiations may be as much an acknowledgment of a process of feature construction as a statement about the distributional properties of the world. At a minimum, however, we are arguing that constructing distinctive shared features at a perceptually specific level is easier than at a more abstract level. More precisely, constructing features shared with at least a few other individuals but only within a category is easier at a more specific level. Regardless of origin, however, specific instantiated features, when present, are more likely to be predictive of categories than are the more abstract informational features.

For the purpose of framing efficient characterizations of categories, however, features at the level of instantiations are inappropriate. Despite its high predictiveness, a particular instantiation covers so small a range of items that a great many such features are required to cover the tremendous variety within most concrete categories. Informational features, selected to be sufficiently abstract to cover most of the manifestations within a category, pay a price in also applying to features in other categories. At this informational level, but only at this level, we have the problem pointed out by Wittgenstein (1953): No single abstract feature is necessary for a category, and no small combination of features is sufficient. The informational feature is excellent for efficient communication, for broad transfer, and for deliberate control of inspection, whereas the instantiation is the more informative for identifying a particular case.

This contrast between the scope and the validity of features is fundamental to their roles in categorization. Instantiated features are selected for their ability, when present, to accurately predict the relevant category. This level is learned in the inspection of new members of a category (strongly facilitated by the experimenters’ selection and manipulation of materials) and is heavily relied on in the identification of new items. In contrast, informational features are selected for their scope and their applicability to a wide range of instantiations. This level is the content of normal communications and, we suspect, a component of the deliberate control of attention. If the materials in an experiment do not represent both levels of features, they are restricting the phenomena that can be explored. If a model does not represent both levels of features, it cannot account for important phenomena in categorization and concept learning.

The Form of Everyday Rules

We have characterized everyday rules (weak rules) as consisting of a list of (informational) features without a decision component, like those produced by our participants in Experiment 1 and as found in many areas of expertise such as medicine. As just argued, we expect most of the features mentioned to be at a relatively abstract, informational level given the communicative intent of stating a rule. However, why should there be no decision procedure? There are several background reasons that are obvious and clearly relevant. However, there is one additional reason that follows from the general theme of this article that should also be considered.

One background possibility for the lack of an explicit decision rule is that people generally have a hard time expressing quantitative information. Possibly a more general reason for not closely specifying weights is that unit weights are sufficient to capture most of the predictive power for new items. As Dawes (1979) and Dawes and Corrigan (1974) have demonstrated, elaborate estimation of weights often improves prediction mainly for the original data set but relatively little for generalization data. Simply checking off the features for and against a hypothesis may capture most of the generalizable variance. However, if this checking off were a normal procedure, we would expect people to offer counting rules quite frequently: rules such as "three out of four" or "most of the features."

An explanation for the usual lack of a decision procedure, consistent with the theme of this article, is that weighting of informational features across items is not the only and maybe not the most useful information. Consider the generalization items in the right column of Figure 2. It might be tempting to call the top item a ramus because of the legs (shared with the prototype ramus), the second item a ramus because of the head (shared with the prototype ramus), and the third and fourth items ramuses because of the pattern and torso shape, respectively (also shared with the prototype). That is, the most useful feature could change from item to item depending on the particular instantiation of the feature. Considering everyday stimuli, different features may be clear under different viewing conditions, and different features may have been subject to distortion because of the vicissitudes of life. It is when feature instantiations are treated as a general class, when the decision is being made on the level of informational features, that the averaging involved in weighting terms across items is the most useful.

In sum, the weak rule form of everyday communications may reflect important aspects of identification procedures. If identification is commonly influenced by familiar feature instantiations that vary from trial to trial, then expressing a weighting of informational features across items may be inappropriate. If the learner's task is to become familiar with a varied array of instantiations, then the most useful information to be given may be the names of the features to be learned about. A list of informational features, then, would express most of what is easy and useful to express.

Two Levels of Grounding of Terms

The first all-too-familiar step for medical students, when faced with a novel medical rule such as "pruritic, violaceous, polygonal-shaped papules," is to look up the terms in a medical dictionary. Even after translation into "itchy, violet, angular bumps," these terms have to be useful for perceptual expectations in order for the rule to guide inspection of the case. By some direct or indirect processes, the terms have to be grounded in visual perception for them to function for a new learner (Glenberg & Robertson, 2000). However, the process described in the introduction and throughout the experiments suggests that this grounding changes rapidly with experience of a particular category. Experience with a number of birds provides a set of instantiations of flying that gradually becomes birdlike flying; experience with cases of lichen planus becomes the basis of lichen planuslike polygonal bumps. The original general language terms *flying* and *polygonal papules* become restricted in scope to no longer apply to manifestations that are inappropriate to the category. Such a set of instantiations becomes exactly the radial, perceptual symbol system described by Barsalou (1999) and Goldstone and Barsalou (1998); the contrast between informational and instantiated features is parallel to the distinction between global and local features made by Solomon and Barsalou (2001). What we hope to add is the linking of this distinction to a somewhat different task analysis and an emphasis on the contrast in operation of the two different referential levels of features.

Unless a term in a rule is grounded at the level of general language, it will not be useful to a beginner. Unless the term becomes restricted to concept-specific usage, it will not reflect experience in a complex world. With experience, the term becomes the name for a set of acceptable instantiations. Our suggestion, then, is that terms have to be grounded at both the general-language and the concept-specific levels to function usefully in concept identification.

Modeling Instantiated Features

In principle, the concept-specific representation of an informational feature could be some sort of an average of the perceptual properties of instantiations previously experienced in that category. A new feature would activate this instantiated representation proportionately to the similarity of its properties to those of the averaged representation. Alternatively, a region of feature space in which all features are assigned to a single category could be characterized as lying within a set of decision boundaries (e.g., see review by Ashby & Maddox, 2005). For categories in which there are smooth transitions among the feature manifestations assigned to the same category (e.g., continuous sensory dimensions such as wavelength or location), either of these could be a reasonable representation. However, as previously mentioned, human legs are as diverse as a baby's legs, a sumo wrestler's legs, a young woman's legs, a hairy old man's legs, and the legs of a victim of starvation. Despite this variety, it is still possible to see legs that would more nearly suggest a bird, a filmmaker's android, or just a bad drawing. That is, similarity to specific subvarieties as well as different from all are important operations in people's judgments within the category of human legs (see Mewhort & Johns, 2005, for a model of such a process applied to recognition memory) and therefore in using these features for a decision about the category of human. This is analogous to the major argument for instances over prototypes as a representation of concepts. In a sense, we are proposing instances on a feature level: a specific version of treating features as categories in themselves, as suggested by Schyns and colleagues (Schyns et al., 1998; Schyns & Murphy, 1994; Schyns & Rodet, 1997).

However, there is a crucial limit to the instances on a feature-level analogy for instantiated features. The operation of instantiated features suggested in this article often is clearly in the service of analytic processing. The instantiated features invoked by a rule are attended to, at least initially, because they are held to be especially relevant properties of the case being categorized. In contrast, one of the important properties of instance models is that they can operate in a more nonanalytic manner. Saying that a current item is similar to an item previously experienced does not require that one have a strong theory about the key elements for the judgment being made. If the two items match on enough properties, the key properties are probably matched also, and the odds are that putting them in the same category is appropriate. Such global matching is clearly not what is happening in most of the experiments in this article, most obviously in Experiment 3 in which the presented rule strongly influenced which instantiations became especially influential. Instantiated features are a way in which heterogeneous perceptual specificity can affect the application of rules and causal explanations.

Obviously, the instantiated features suggested in this article could also operate as components of whole instances. The results of prior rule-directed attention could be embodied in instance representations by greater weighting of the attended features. However, the perceptual specificity of instantiated features potentially allows this weighting to be done in a more flexible manner than by weighting whole dimensions across all items, as is done in the generalized context model (Nosofsky, 1984) and ALCOVE (Kruschke, 1992). Clearly, the results with the interfering-features items in Experiment 1 (see Figure 2, right column) imply weighting that is not done a whole dimension at a time. What is not a matter of conjecture, however, is that there are effects of perceptual specificity in these experiments that cannot be accounted for by instantiated features. Experiment 6 demonstrated that learners responded to the joint familiarity of combinations of instantiated features and not just to a sum of the familiarity of the individual feature instantiations. That is, some relational information was learned and used in the plausibility ranking given by the participants.

It is important, however, to remember that the conditions in these experiments clearly favor the more direct effect of individual features rather than the indirect contribution through selecting prior whole instances. The difference between the all-novel conditions and the perceptual-interference conditions of Experiments 1–4 was the familiarity of a single instantiated feature, not the overall similarity of whole instances. This was meant to model the situation in which a beginner's eye is caught by a particularly meaningful feature, a common report in medical learning. Second, the structure of the learning materials contributed to focusing on features rather than instances. Because the items within a category shared many instantiated features, the individual items were less distinctive than if each item had many unique instantiations (e.g., Allen & Brooks, 1991; Regehr & Brooks, 1993). Further, all items had the same number of parts, and both categories had an identical dimensional structure. This regular structure makes it easy for a participant to compare items for the parts that contrast or are shared, but it is not a characteristic of natural categories such as medical cases. Finally, of course, several of these experiments gave the participants weak rules that were specified in terms of parts. We think that these are interesting conditions to investigate, especially for formal teaching, but they certainly are not neutral for what is learned.

In modeling the phenomena described in this article, we suspect that we will ultimately have to deal with instantiations and informational features either in different systems or with different specification of their processing contexts (e.g., Ashby & Maddox, 2005). However, we believe that it is important to first explore the explanatory power of representing features with a wide scope but moderate predictiveness in contrast to features with a narrow scope but very high predictiveness. A simple, two-layer connectionist model in which the instantiated–informational distinction is modeled solely as features with these contrasting scope–predictiveness properties is able to capture the main phenomena in this article (Hannah, Jankowitz, & Brooks, 2004). It also produces the associative blocking effects characteristic of several connectionist models, including Kruschke's ADIT and ATRIUM (Kruschke, 2001), that also occur in our materials. The relations between instantiated and informational features cannot be arbitrary, but the relative weighting of them is easily changed by experimental manipulations such as number of instantiations per informational feature and the cue validity of the informational features. Also, as suggested by the results of Experiment 6, a more general model will ultimately have to include coordination of feature-based and instance-based categorization, similar to that provided in ATRIUM by Erickson and Kruschke (1998) and Kruschke (2001).

Conclusion

Finally, let us return to the undergraduate's definition of a bird: "A bird sings, flies, and has feathers." The terms are grounded at a general-language level useful for a beginner or they would not be acceptable as an answer to an everyday "How do you know?" question. However, they clearly have a concept-specific extension because the student did not accept the example of a diva with a feather boa on a jet as relevant to the rule. The decision component was missing the student's rule at least partly because the weighting of particular characteristics is not stable across perceptual presentations or novel items. What is given is useful as long as people accept it as instructions for perceptual learning, not as an identification procedure. These formally inadequate everyday rules say much of what can be said within the efficiency limits of acceptable communication. The same seems to be true of the identification rules given in areas of formal instruction, despite the obviously greater care given to their formulation. These contrasts between the knowledge used in general communication or personal attention control and that used in specific identification seem central to understanding adults' category learning.

References

Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General, 120,* 3–19.

American Psychiatric Association. (1994). *Diagnostic and statistical manual of mental disorders* (4th ed.). Washington, DC: Author.

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology, 56,* 149–178.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22,* 577–609.

Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). The role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General, 120,* 278–287.

Chin-Parker, S., & Ross, B. H. (2004). Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 216–226.

Chuang, T.-Y., & Stitle, L. (n.d.). *Lichen planus*. Retrieved April 11, 2006, from http://www.emedicine.com/DERM/topic233.htm

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34,* 571–582.

Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin, 81,* 95–106.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140.

Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition, 65,* 263–297.

Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language, 43,* 379–401.

Goldstone, R. L., & Barsalou, L. (1998). Reuniting perception and conception. *Cognition, 65,* 231–262.

Hannah, S. D., & Brooks, L. R. (2006). *Specificity of feature representation influences categorization decision processes*. Manuscript submitted for publication.

Hannah, S. D., & Brooks, L. R. (in press). Producing biased diagnosis with unambiguous stimuli: The role of feature instantiations. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Hannah, S. D., Jankowitz, D., & Brooks, L. R. (2004, June). *Modeling the effect of feature familiarity in categorization*. Paper presented at the 14th Annual Meeting of the Canadian Society for Brain, Behaviour and Cognitive Science, St. John's, Newfoundland and Labrador, Canada.

Kaplan, A. S., & Murphy, G. L. (2000). Category learning with minimal prior knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 829–846.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology, 45,* 812–863.

LeBlanc, V. R., Norman, G. R., & Brooks, L. R. (2001). Effect of a diagnostic suggestion on diagnostic accuracy and identification of clinical features. *Academic Medicine, 76,* S18–S20.

Malt, B. C., & Smith, E. E. (1983). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior, 23,* 250–269.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431–467.

Markman, A. B., & Maddox, W. T. (2003). Classification of exemplars with single and multiple feature manifestations: The effects of relevant dimension variation and category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 107–117.

Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin, 129,* 592–615.

Mewhort, D. J. K., & Johns, E. E. (2005). Sharpening the echo: An iterative-resonance model for short-term recognition memory. *Memory, 13,* 300–307.

Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 904–919.

Murphy, G. L., & Wisniewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications* (pp. 23–45). Chichester, England: Ellis Horwood.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 104–114.

Regehr, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General, 122,* 92–114.

Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 1141–1159.

Rehder, B., & Ross, B. H. (2001). Abstract coherent categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 1261–1275.

Schyns, P. G., Goldstone, R. L., & Thibaut, J. (1998). Development of features in object concepts. *Behavioral and Brain Sciences, 21,* 1–54.

Schyns, P. G., & Murphy, G. L. (1994). The ontogeny of part representation in object concepts. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 31, pp. 305–349). New York: Academic Press.

Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 681–696.

Sloman, S. A., & Rips, L. J. (1998). Similarity as an explanatory concept. *Cognition, 65,* 87–101.

Solomon, K. O., & Barsalou, L. W. (2001). Representing properties locally. *Cognitive Psychology, 43,* 129–169.

Spalding, T. L., & Murphy, G. L. (1999). What is learned in knowledge-related categories? Evidence from typicality and feature-frequency judgments. *Memory & Cognition, 27,* 856–867.

Wattenmaker, W. D. (1993). Incidental concept learning, feature frequency, and correlated properties. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 203–222.

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. (2000). Learning categories composed of varying instances: The effect of classification, inference, and structural alignment. *Memory & Cognition, 28,* 64–78.

Appendix

Feature List Display Instructions for Experiments 2 and 3

Experiment 2

"Here is a list of features that may help you: Bleebs usually have two legs, stripes, rounded head and rounded torso, while a Ramus usually has four legs, dots, angular head, and an angular torso."

Experimenter displays feature list and Pairs 1–10 with names beneath (approximately 10 s each).

Experiment 3

"Here is a list of features that may help you."

Experimenter displays feature list on overhead and reads either Feature Set A or Feature Set B.

Feature Set A. "Bleebs usually have a round, whistle-shaped torso or body; two plump, stationary feet and legs; and a short, wide, fanlike crest. Ramuses usually have a six-sided torso or body; four thin-stepping feet or legs; and a tall, treelike crest."

Feature Set B. "Bleebs usually have an egg-shaped head on a short, wide neck; dark, widely spaced stripes; and a short, feathery tail. Ramuses usually have a four-sided head on a long, thin neck; scattered dots on their body; and a long curved tail."

Experimenter shows first pair and points out that "Indeed, this bleeb has [feature list], and this ramus has [feature list]." Experimenter then displays feature list and Pairs 1–15 with names beneath (approximately 10 s each), announcing bleeb and ramus for each displayed pair.

Received March 3, 2005
Revision received July 21, 2005
Accepted August 2, 2005 ■