

Noncategorical approaches to feature prediction with uncertain categories

Christopher Papadopoulos · Brett K. Hayes · Ben R. Newell

© The Psychonomic Society 2010

Abstract In four experiments, we investigated how people make feature predictions about objects whose category membership is uncertain. Artificial visual categories were presented and remained in view while a novel instance with a known feature, but uncertain category membership was presented. All four experiments showed that feature predictions about the test instance were most often based on feature correlations (referred to as *feature conjunction reasoning*). Experiment 1 showed that feature conjunction reasoning was generally preferred to category-based induction in a feature prediction task. Experiment 2 showed that people used all available exemplars to make feature conjunction predictions. Experiments 3 and 4 showed that the preference for predictions based on feature conjunction persisted even when category-level information was made more salient and inferences involving a larger number of categories were required. Little evidence of reasoning based on the consideration of multiple categories (e.g., Anderson, (Psychological Review, 98:409–429, 1991)) or the single, most probable category (e.g., Murphy & Ross, (Cognitive Psychology, 27:148–193, 1994)) was found.

Keywords Induction · Category uncertainty · Feature correlations

One of the most important functions of categories is to allow people to make inductive predictions about novel

instances. Much of the existing research on inductive reasoning has involved cases in which the category membership of a target instance is known with certainty (see Heit, 2000, for a review). However, given our inherently complex and uncertain world, it seems likely that we will sometimes have to make inductive predictions about instances before their category membership has been determined. For example, a patient visiting a clinical psychologist may present with auditory and visual hallucinations. This symptom profile might suggest schizophrenia as a likely diagnosis, but other diagnoses (e.g., substance abuse, bipolar disorder) are possible. Before a certain diagnosis can be made, however, the clinician may want to make inductive predictions about the patient (e.g., “Is she likely to harm herself?”). The central question that we address is how people make such predictions when category membership is uncertain.

According to Anderson’s (1991) Rational model, people make predictions under category uncertainty by considering all of the categories to which the novel instance may belong, a process deemed *multiple-category reasoning*. In the example above, this would involve (1) computing the relative probability of various diagnostic categories, given the presenting symptoms; (2) computing the conditional probabilities of a predicted feature (e.g., likelihood of self-harm) for each category; and (3) integrating this information over the various category alternatives, using Bayes’ rule.

Although considered a normative approach, there is little evidence that people reason in this way. Instead, most evidence suggests that people rarely consider multiple categories in feature induction with uncertain categories (see Murphy & Ross, 2007, for a review). In most cases, people make such predictions on the basis of information drawn only from the category to which an instance is most likely to belong (schizophrenia, in the clinical example).

C. Papadopoulos (✉) · B. K. Hayes · B. R. Newell
School of Psychology, University of New South Wales,
Sydney, NSW 2052, Australia
e-mail: c.papadopoulos@unsw.edu.au

B. K. Hayes
e-mail: B.Hayes@unsw.edu.au

B. R. Newell
e-mail: Ben.Newell@unsw.edu.au

Murphy and Ross (1994) found evidence for such *single-category reasoning* in a paradigm in which participants were shown geometric figures, varying in shape and color, which were said to be drawn by one of a number of different children. At test, a novel instance was presented that had a given feature (e.g., square) but whose category membership was uncertain (because more than one child drew squares). Participants were asked to choose the target category for the novel instance (i.e., which child was most likely to have drawn it), predict the other feature (e.g., color) that it was most likely to have, and rate the probability of this predicted feature. For some inductive problems, consideration of nontarget categories altered the conditional probability of the predicted feature. Even though people explicitly acknowledged that category membership of the test item was uncertain, their probability ratings showed that they ignored this uncertainty when making feature predictions, basing these on information from the target category alone. This finding has since been widely replicated with both artificial (e.g., Verde, Murphy, & Ross, 2005) and natural (e.g., Murphy & Ross, 2005; Ross & Murphy, 1996) categories.

Feature conjunction reasoning

The finding that people generally prefer reasoning on the basis of single, as opposed to multiple, categories when category membership is uncertain is important. However, we argue that the focus on single- and multiple-category reasoning has led to the neglect of an alternative reasoning strategy that may actually be more common in cases of induction with uncertain categories. This approach, which we refer to as *feature conjunction* reasoning, involves making feature predictions on the basis of a comparison between the features of a given test instance and similar exemplars. Although it is possible that people may restrict this comparison to exemplars within the target category (see Murphy & Ross, 2010a, and Experiment 2), we believe that feature conjunction predictions are usually based on a comparison with all available exemplars that have features in common with the target, ignoring category boundaries. In the clinical example, feature conjunction would involve making a prediction by comparing the new patient with previously experienced patients with the same or a similar symptom profile, regardless of the explicit diagnosis.

This approach and its relationship to single-category and multiple-category reasoning can be illustrated in more detail using Fig. 1, which shows two categories of drawings attributed to different artists. When presented with a novel exemplar of a particular shape (e.g., a square), most people should have little difficulty in identifying the target category (i.e., Peter, since he drew the most squares) but

should also recognize that category membership is uncertain, because both categories contain squares. If asked to predict the value of the second feature dimension in the test item, a multiple-category approach (Anderson, 1991) would lead to the prediction “purple,” since this color has a higher relative frequency when features from both categories are considered. A single-category approach (Murphy & Ross, 1994), on the other hand, would restrict consideration of feature frequency to the target category alone, leading to a prediction of “aqua,” the most common color in this category.

Note that both the multiple-category and single-category approaches assume that people treat the given feature (e.g., shape) and the predicted feature (e.g., color) as conditionally independent. In other words, the given feature is used to identify only the most likely categories (or category). Relationships between features play no further part in generating feature predictions, which are based on the distribution of the to-be-predicted feature across all exemplars within the relevant categories. In contrast, the feature conjunction approach assumes that people make predictions by examining exemplars that have features in common with the test item. In the Fig. 1 example, squares in both categories would be examined, and the color that most frequently occurs together with squareness (i.e., “red”) would be predicted. In this case, people are using their knowledge of relationships between object features, rather than category membership, to make a feature prediction.

There are good reasons for expecting that, given the opportunity, people will use feature conjunction to make predictions about the unknown features of specific objects. One reason is that, contrary to the assumptions of single-category and multiple-category approaches, people often do not treat features as conditionally independent within categories. Although some studies have shown that people do not notice within-category feature correlations when classifying objects (e.g., Murphy & Wisniewski, 1989), others have shown that a sensitivity to such correlations does emerge with extended learning (Thomas, 1998). More important, people seem sensitive to within-category feature correlations between features when asked to make predictions about the missing features of category exemplars (Chin-Parker & Ross, 2002; Murphy & Ross, 1994, Experiments 7, 8, and 11; Murphy & Ross, 2010a; Yamauchi, Love, & Markman, 2002).

A second reason for expecting that people may use exemplar similarity and correlated features as a basis for reasoning is the body of work showing that specific similarity to familiar exemplars affects categorization. Even when there are clear rules for categorizing, classification decisions are influenced by the similarity of exemplars from the same or other categories (e.g., Allen & Brooks, 1991; Regehr & Brooks, 1993; Thibaut & Gelaes, 2006). When

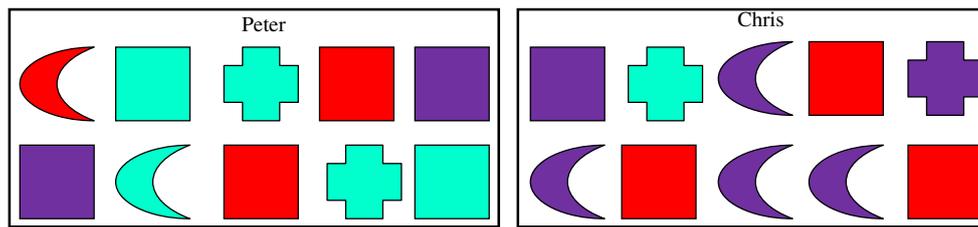


Fig. 1 Example of shape/color categories used in Experiment 1. The critical question was “I have a drawing of a square. What color do you think it is?” Single-category prediction = “aqua”; multiple-category prediction = “purple”; feature conjunction prediction = “red.”

category membership is uncertain, exemplar similarity may play an even greater role, perhaps extending to inductive predictions about object features.

If feature conjunction is as prevalent as we suppose, it is reasonable to ask why little evidence of this approach has been reported in previous work on feature prediction under uncertainty. The first answer is that in many previous studies, feature conjunction has been explicitly excluded from the experimental design. Murphy and Ross (1994, Experiments 5 and 6), minimized the likelihood that people would use feature conjunction by making sure that no feature was paired with the given feature more frequently than were any other within or between categories (so feature conjunction could not yield a clear feature prediction).

More seriously, in other cases, the predictions based on feature conjunction have been confounded with those of category-based approaches. In Murphy and Ross (1994, Experiments 1–3), for example, a feature conjunction approach led to the same feature predictions as the single-category approach, the latter being the default strategy, according to the authors. The finding of Verde et al. (2005) that speeded induction decisions were influenced by information from outside of the target category was interpreted as supporting multiple-category reasoning. This finding, however, is equally compatible with an explanation based on a feature conjunction approach that ignores category boundaries (Newell, Paton, Hayes, & Griffiths, 2010).

Hayes, Ruthven, and Newell (2007) found some evidence that people use feature conjunction reasoning when making feature predictions under category uncertainty, even when alternative category-based approaches were available. Participants were presented with artificial category structures and were asked to make a prediction about the probability of target instances having some critical feature, given that they had some other feature. Feature base rates were manipulated so that feature likelihood ratings would increase from the baseline to the experimental condition if people used feature conjunction reasoning but would remain relatively constant if they used category-based reasoning. The data suggested that at least 50% of participants used feature conjunction reasoning. One limitation of this study, however, is that it was not clear how the

remaining participants were approaching induction. The finding of no difference between the baseline and experimental conditions does not tell us what type of reasoning strategy these participants were following.

The main aim of the present experiments was to further examine the degree to which people use feature conjunction, as compared with multiple-category or single-category, reasoning when making feature predictions with uncertain categories. Participants were shown a series of induction problems involving category uncertainty. Each problem consisted of categories with a distribution of features like that in Fig. 1. After studying these categories, participants were shown a novel test item whose category membership was uncertain (e.g., a square in the Fig. 1 example) and were asked to make a feature prediction (e.g., it’s most likely color). As in previous work (e.g., Murphy & Ross, 1994), we employed a *decision-only* paradigm whereby the relevant categories remained visible while test predictions were made. The rationale was that by presenting all relevant category information, the prediction process could be studied independently of learning and memory processes.

Critically, each induction problem was designed so that the reasoning strategies of interest (feature conjunction, single-category, and multiple-category reasoning) led to divergent feature predictions (see Table 1 for details). This meant that we were able to establish whether a participant was using feature conjunction or a type of category-based reasoning for each induction problem.¹ This paradigm was used to compare the relative prevalence of feature conjunction with single-category and multiple-category reasoning (Experiment 1). Subsequent experiments compared categorical and noncategorical versions of feature conjunction reasoning (Experiment 2) and examined whether use of feature conjunction reasoning depends on category salience (Experiment 3). Experiment 4 examined whether the

¹ Previous work (e.g., Murphy & Ross, 2010b) has employed a similar paradigm whereby different reasoning strategies made divergent predictions. However, this work did not examine predictions based on feature conjunction.

Table 1 Experiment 1: Example of probabilities of feature choices on the induction test based on each induction strategy for Fig. 1 categories and a test exemplar with given feature “square” (Peter is the target category, Chris is the nontarget category). The most likely feature choice for each strategy is given in bold

Feature alternatives	Single category (Murphy & Ross, 1994)	Multiple category (Anderson, 1991)	Feature conjunction
$p(\text{respond } \textit{aqua})$	$= p(\text{aqua} \text{Peter})$ $= 5/10$ $= \mathbf{0.5}$	$= [p(\text{Peter} \text{square}) * p(\text{aqua} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{aqua} \text{Chris})]$ $= [0.6 * 0.5] + [0.4 * 0.1]$ $= 0.34$	$= p(\text{aqua} \text{square})$ $= 2/10$ $= 0.2$
$p(\text{respond } \textit{purple})$	$= p(\text{purple} \text{Peter})$ $= 2/10$ $= 0.2$	$= [p(\text{Peter} \text{square}) * p(\text{purple} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{purple} \text{Chris})]$ $= [0.6 * 0.2] + [0.4 * 0.6]$ $= \mathbf{0.36}$	$= p(\text{purple} \text{square})$ $= 3/10$ $= 0.3$
$p(\text{respond } \textit{red})$	$= p(\text{red} \text{Peter})$ $= 3/10$ $= 0.3$	$= [p(\text{Peter} \text{square}) * p(\text{red} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{red} \text{Chris})]$ $= [0.6 * 0.3] + [0.4 * 0.3]$ $= 0.3$	$= p(\text{red} \text{square})$ $= 5/10$ $= \mathbf{0.5}$

dominance of feature conjunction reasoning would generalize to induction involving more than two categories.

Experiment 1

The aim of this experiment was to compare feature conjunction with single-category and multiple-category approaches to inductive reasoning when the category membership of a test instance was uncertain. Unlike in many previous studies (e.g., Murphy & Ross, 1994), the stimuli were designed so that any of these approaches could be used to make a feature prediction for each test instance. Critically, however, the three forms of reasoning led to qualitatively different predictions, so that there was a clear index of the extent to which participants relied on each approach. On the basis of the results of Hayes et al. (2007), it was expected that a large proportion of participants would employ feature conjunction reasoning, rather than category-based approaches.

Method

Participants

Twenty-five undergraduate students from the University of New South Wales, Australia, participated for course credit in an introductory psychology course ($M = 20.48$ years).

Materials

Eight uncertain induction problems were developed. Each problem consisted of two categories of drawings (i.e. a target and a nontarget category) made up of 10 exemplars

each and a novel test exemplar with a given feature on one dimension, but with the feature value on the other dimension unspecified. The categories used for each problem had the same distribution of features as the categories in Fig. 1. However, the appearance of category pairs differed across problems. For four problems, category exemplars varied in their color and shape, each of which could take three feature values. For these problems, participants were told that the categories were pictures drawn by different children. The other four problems used exemplars that varied in terms of more complex irregular shapes filled with black-and white-patterns or colored textures (see the Appendix, Fig. A1). For these problems, participants were told that the categories were drawings by different students in a graphic design course. These graphic design stimuli were used to increase the total number of problems presented to each participant and, hence, obtain more reliable estimates of individual patterns of reasoning. The relevant categories for each problem were presented on a laminated A4 page in portrait orientation.

Procedure

All the participants were presented with all eight uncertain induction problems. At the beginning of each problem, the two categories were shown, and participants were told that these represented a sample of the typical drawings done by the child (or graphic artist) whose name appeared above the category (as in Fig. 1). The relative position of the target and nontarget categories at the top and bottom of the A4 page was counterbalanced across problems. Participants were given 1 min to study the categories. After that time, the categories remained in view, but a novel test exemplar was presented. This exemplar had a given feature on one dimension (e.g., shape), but the feature value on the other

dimension (e.g., color) was not presented. When the test exemplar was a shape, it was depicted as a black-and-white outline. When the test exemplar was a color, it was depicted as a color swatch in an irregular shape. Participants were told that this was a “new picture that may have been drawn by one of the children (or graphic design students)” and that they would have to answer a number of questions about it (printed on a separate A4 sheet).

The first questions involved identifying the target category of the test exemplar (e.g. “I have a picture of a square. Which child do you think is most likely to have drawn it?”). This allowed us to check that participants were able to correctly identify the target category, a crucial assumption underlying the predictions of category-based strategies. The second question involved rating confidence in this judgment on a scale from 0 (*not at all confident*) to 100 (*completely confident*). This allowed us to check whether participants perceived the category membership of the test exemplar as uncertain. Questions 3 and 4 were fillers and involved counting up feature values on the given dimension (e.g., “How many crosses did Peter draw?”). The final two questions involved the key inductive predictions. Participants were asked to indicate what other feature value they thought was most likely to be found in the novel test exemplar along with the given feature (e.g. “I have a drawing of a square. What color do you think it is?”). The three possible feature values were depicted perceptually, using shape drawings with a black outline (for color-to-shape predictions) or color swatches with irregular shapes (for shape-to-color predictions). Participants circled the response option they thought was correct and then rated their confidence in this judgment on a 0–100 scale.

All the problems were designed so that feature conjunction, multiple-category, and single-category approaches gave rise to different predictions about which feature would be chosen. The details of the feature predictions, in terms of the relative probabilities of each feature based on the assumptions of each reasoning approach, are given in Table 1. Note that each approach predicts that a different feature is the most likely one to be chosen on the induction test question. However, in this case, the multiple-category approach predicts similar probabilities for two features (aqua and purple), meaning that predictions based on this approach may be hard to distinguish from those based on the single-category approach. This problem was addressed in Experiment 4.

For half the problems, people were given the shape of a test instance and had to predict its color/fill pattern, whereas for the remaining problems, these roles were reversed. Furthermore, the distribution of feature values within the predicted dimension was counterbalanced across participants to control for differences in the perceived salience of feature values. For example, the color “red” was the feature

predicted by multiple-category, single-category, and feature conjunction approaches on an equal number of occasions.

After all the questions for an induction problem were completed, it was removed, and the next problem was presented. Presentation of the four “children’s drawings” and the four “graphic design” problems was blocked, with block order counterbalanced across participants.

Results

Preliminary analyses indicated that feature predictions did not vary across the “children’s drawings” and “graphic design” problems or across counterbalanced versions of the task. Hence, all the subsequent analyses are collapsed across these factors.

Participants were extremely accurate in identifying the target category ($M = .99$ correct). Nevertheless, they gave modest confidence ratings for these categorization judgments ($M = 64.85$, $SD = 10.79$), suggesting that they understood that category membership of test instances was uncertain.

The most important analyses pertain to feature predictions about the test instance when the given feature was presented. As in past studies (e.g., Murphy & Ross, 1994), feature prediction responses were included in the analysis only for problems where the target category was identified correctly. For the feature prediction task, the mean proportion of responses that were consistent with each of the three reasoning approaches was calculated for each participant. An overwhelming proportion of the feature predictions was consistent with the feature conjunction approach ($M = .97$, $SD = .10$), with the mean proportion well above a chance value of .33, $t(24) = 30.67$, $p < .001$. The proportion of feature predictions consistent with single-category ($M = .015$, $SD = .08$) and multiple-category ($M = .015$, $SD = .04$) reasoning was close to floor. The mean confidence rating for feature predictions was reasonably high ($M = 74.63$, $SD = 18.69$) and did not vary across reasoning strategies.

Individual-participant profiles were also calculated in order to determine the extent to which participants displayed “consistent” use of a particular reasoning strategy (defined as at least five out of eight feature predictions based on the same reasoning strategy). All but 1 participant consistently used feature conjunction reasoning. The remaining participant did not show consistent use of any strategy.

Discussion

This experiment showed that when people made inductive predictions about instances whose category membership was uncertain, they showed an overwhelming preference

for predictions based on feature conjunction. There was very little evidence of category-based reasoning. Most participants showed no respect for category boundaries when making feature predictions, appearing to base their predictions on a comparison with all the exemplars that were similar to the test instance on the given feature dimension. Feature predictions were based on a consideration of the frequency of feature co-occurrence within this set of exemplars. This represents a radical departure from the findings of previous work, which has assumed that people treat given and predicted features as conditionally independent, with people making feature predictions on the basis of the overall frequency of features in all of the categories to which the target instance may belong (Anderson, 1991) or the overall frequency in the target category alone (e.g., Murphy & Ross, 1994). These previous studies may have seriously underestimated the prevalence of feature conjunction by eliminating it as a viable prediction option or confounding it with category-based predictions.

Experiment 2

Experiment 1 showed that when people have to make feature predictions with uncertain categories, they show an overwhelming preference for feature conjunction. It is important to note, however, that feature conjunction, as operationalized in **Experiment 1**, differed in at least two important respects from previous category-based approaches. First, unlike in the single-category approach (e.g., Murphy & Ross, 1994), it was assumed that feature conjunction involved examination of the entire set of exemplars that matched the test instance on a given feature. Second, unlike in the multiple category approach (e.g., Anderson, 1991), it was assumed that people do not treat feature dimensions as conditionally independent. This assumption seems well justified, given previous work suggesting that people are sensitive to feature co-occurrences within categories (e.g., Chin-Parker & Ross, 2002; Murphy & Ross, 2010a).

Experiment 1 did not differentiate between these two aspects of feature conjunction. The aim of the present experiment was to disentangle the issue of feature independence from consideration of multiple exemplars by comparing the use of two kinds of feature conjunction in induction with uncertain categories. Both kinds of conjunctions assume that feature dimensions are treated as nonindependent. Multiple feature conjunction generates predictions based on examination of all the available exemplars, regardless of category boundaries (as in **Experiment 1**). In contrast, single-category feature conjunction involves only the examination of exemplars within the

target category. In this experiment, we designed a stimulus set where qualitatively different feature predictions could be generated by these two versions of feature conjunction.

If the results of **Experiment 1** reflect mainly the fact that people do not treat feature dimensions as independent, there should be a substantial proportion of predictions consistent with each version of feature conjunction. Given the bias toward focusing on the information within the target category found in previous studies (e.g., Murphy & Ross, 2005, 2010a), it is also possible that in a head-to-head comparison, people may favor the single-category version of feature conjunction. If, however, people usually consult all the available exemplars, there should be a preference for multiple-feature conjunction.

Method

Participants

Twenty-four undergraduate students participated for course credit in an introductory psychology course ($M = 18.83$ years).

Design and procedure

The general design and procedure followed those in **Experiment 1**. In this case, however, induction problems were designed to compare predictions based on multiple-category feature conjunction and single-category feature conjunction, precluding the use of either of the category-based strategies examined in the previous experiment. Eight induction problems were developed and administered in the same way as in **Experiment 1**.

An example of the categories used is given in Fig. 2. In this example, the given feature was “square,” the target category was Peter, and the nontarget category was Chris. Multiple-feature conjunction would involve examining all of the squares in the target (Peter) and nontarget (Chris) categories and noticing that “red” was the most common color for squares. Single-category feature conjunction would involve attending only to the subset of squares that appear in the target category (Peter) and noticing that they were most likely to be “aqua.” Note that a single-category strategy (Murphy & Ross, 1994) would produce an ambiguous feature prediction, since Peter drew an equal number of aqua and red shapes. Similarly, a multiple-category approach (Anderson, 1991) would also lead to an ambiguous feature prediction (see Table 2 for probabilities associated with the feature predictions). Furthermore, note that the overall base rates of the critical response options based on the two conjunction strategies (e.g., red vs. aqua in Fig. 2) were equated for all induction problems.

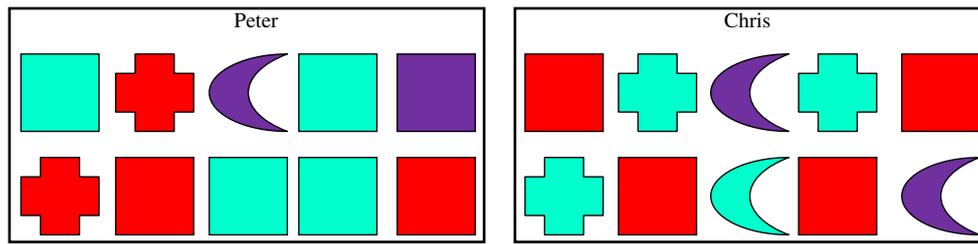


Fig. 2 Example of shape/color categories used in Experiment 2. The critical question was “I have a drawing of a square. What color do you think it is?” Single-category feature conjunction prediction = “aqua”; multiple-category feature conjunction prediction = “red.”

Results and discussion

Preliminary analyses indicated that feature predictions did not vary across the children’s drawing and graphic design problems or across counterbalanced versions of the task. All the subsequent analyses are collapsed across these factors.

Participants always identified the target category correctly, except for 1 participant on one problem. Again, the mean confidence rating for category judgments was modest ($M = 69.89, SD = 7.80$), suggesting that participants recognized the uncertainty in category membership.

The proportion of predictions consistent with each feature conjunction approach was calculated for each participant. Predictions were included only for problems where the target category was correctly identified. The mean proportion of feature predictions consistent with multiple-feature conjunction was almost at ceiling ($M = .98, SD = 0.06$) and significantly above chance, $t(23) = 39.00, p < .001$, whereas the mean proportion of feature predictions consistent with single-category feature conjunction was close to floor ($M = .02, SD = .06$). Confidence in feature predictions was

moderate ($M = 71.23, SD = 15.97$). Individual induction profiles were examined using the same consistency criterion as that in Experiment 1. Every participant consistently used multiple-feature conjunction.

There was again an overwhelming preference for feature conjunction reasoning based on an examination of exemplars from all the available categories, rather than just from exemplars within the target category. This shows that the preference for feature conjunction over category-based approaches in Experiment 1 was a joint product of participants assuming that values on the feature dimensions were correlated and examining these correlations across category bounds.

The low level of single-category feature conjunction responding in this experiment seems inconsistent with that in Murphy and Ross (2010a), who found evidence that people frequently use within-category correlations, rather than category-based information, to make feature predictions. Notably, however, Murphy and Ross (2010a) did not compare predictions on the basis of feature correlations within categories (single-category feature conjunction reasoning) with predictions based on feature correlations

Table 2 Experiment 2: Example of probabilities of feature choices on the induction test based on each induction strategy for Fig. 2 categories and a test exemplar with given feature “square” (Peter is

the target category, Chris is the nontarget category). The most likely feature choice for each strategy is given in bold

Feature alternatives	Single-category feature conjunction	Multiple feature conjunction	Single category (Murphy & Ross, 1994)	Multiple category (Anderson, 1991)
$p(\text{respond aqua})$	$= p(\text{aqua} \text{Peter square})$ $= 4/7$ $= 0.57$	$= p(\text{aqua} \text{square})$ $= 4/11$ $= 0.36$	$= p(\text{aqua} \text{Peter})$ $= 4/10$ $= 0.4$	$= [p(\text{Peter} \text{square}) * p(\text{aqua} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{aqua} \text{Chris})]$ $= [0.64 * 0.4] + [0.36 * 0.4]$ $= 0.4$
$p(\text{respond purple})$	$= p(\text{purple} \text{Peter square})$ $= 1/7$ $= 0.14$	$= p(\text{purple} \text{square})$ $= 1/11$ $= 0.09$	$= p(\text{purple} \text{Peter})$ $= 2/10$ $= 0.2$	$= [p(\text{Peter} \text{square}) * p(\text{purple} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{purple} \text{Chris})]$ $= [0.64 * 0.2] + [0.36 * 0.2]$ $= 0.2$
$p(\text{respond red})$	$= p(\text{red} \text{Peter square})$ $= 2/7$ $= 0.29$	$= p(\text{red} \text{square})$ $= 6/11$ $= 0.55$	$= p(\text{red} \text{Peter})$ $= 4/10$ $= 0.4$	$= [p(\text{Peter} \text{square}) * p(\text{red} \text{Peter})] + [p(\text{Chris} \text{Square}) * p(\text{red} \text{Chris})]$ $= [0.64 * 0.4] + [0.36 * 0.4]$ $= 0.4$

across categories (multiple-feature conjunction reasoning). In light of the present results, it appears that although people can use within-category correlations for feature prediction, when given the alternative of making predictions on the basis of feature conjunctions that cross category bounds, they show a preference for the latter approach.

Experiment 3

In [Experiment 3](#), we examined the extent to which our earlier findings regarding the preference for feature conjunction reasoning can be generalized to cases in which category-level information is made more prominent. The labels given to the categories in [Experiments 1 and 2](#) (children's/graphic design students' drawings) were arbitrary. In addition, the level of within-category similarity of the exemplars was considerably lower than that found in the family resemblance categories favored in many other studies of category learning and induction (e.g., Yamauchi & Markman, 1998). Past work (e.g., Patalano, Chin-Parker, & Ross, 2006) showed that perceived category coherence is positively associated with the use of category membership as a basis for inductive prediction. It may be that participants saw the categories in [Experiments 1 and 2](#) as ad hoc collections, with category labels (artist names) providing few clues for feature prediction. If this was the case, it is perhaps unsurprising that people used a non-categorical approach to derive feature inferences.

It is important to note that many previous studies using artificial stimuli broadly similar to those in our experiments (e.g., Murphy & Ross, 1994; Verde et al., 2005) have shown that people use these categories as the basis for inductive prediction. Nevertheless, it is possible that feature conjunction reasoning is used in preference to category-based strategies only when the categories in question are perceived as uninformative or lacking coherence.

In [Experiment 3](#), therefore, we reexamined multiple-feature conjunction and category-based reasoning under conditions in which we varied the extent to which category-level information was made salient. Note that, in this context, *salience* refers to the degree to which category-level information was highlighted when the experimental stimuli were presented. The categories used in the low-category-salience and high-category-salience conditions had the same distribution of features as those in [Experiment 1](#). In the high-salience condition, however, the categories were presented as different kinds of viruses, with exemplar features representing different structural parts of individual viruses. A salient perceptual feature ("viral outer shell") was added to the member of each category to enhance category distinctiveness and increase the similarity of

members within each category. Given that categories of living things are often perceived as sharing a range of both known and unknown features (Gelman, 2003), these manipulations were intended to promote the belief that the categories were informative and meaningful groupings. By contrast, exemplars in the low-salience condition were presented with little rationale for category structure and arbitrary category labels. The critical question was whether feature conjunction would remain the dominant strategy when category-level information was made more salient.

Method

Participants

Twenty-four undergraduate students participated for course credit in an introductory psychology course ($M=19.34$ years). Equal numbers were randomly allocated to the low- and high-salience conditions.

Procedure

The salience of category-level information was manipulated between subjects. In the high-salience condition, participants were told that they had to learn about the features of two kinds of recently discovered viruses ("Sirus," "Karplek"). These viruses were said to be composed of two parts. The first was an "outer shell." Members of each category had different outer shells, but all exemplars within a category had an identical shell. Note that this outline could be used to discriminate between members of the respective categories but was not otherwise relevant to the feature predictions required for test instances. The second part of the virus was an "inner body structure," presented inside the viral shell (see [Fig. 3](#) for an example). These body structures were identical in appearance and statistical structure to the exemplars used in the graphic design problems in [Experiment 1](#). Hence, they allowed for the same demarcation between feature predictions based on feature conjunction, single-category, and multiple-category reasoning. After studying the categories relevant to each induction problem for 1 min, a novel test exemplar was presented *without* the viral outline. The test exemplar was presented as a new viral internal body part that could belong to either of the two categories of viruses being presented. Note that the viral outline was never used as a given or predicted feature in induction questions. Participants answered six questions about the test exemplar (including categorization, filler, and feature prediction) analogous to those used in [Experiment 1](#). Each participant completed four problems of this type, with presentation order randomized.

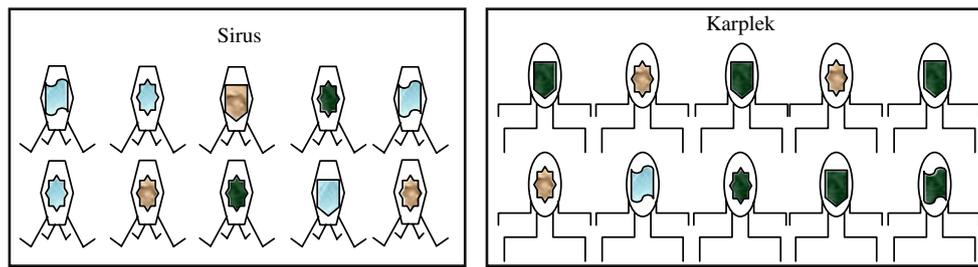


Fig. 3 Example of categories from the high-salience condition in Experiment 3. The critical question was “I have a capsid (an internal viral structure) in the shape of a star. What kind of protein material do

you think it is most likely to be filled with?” Single-category prediction = “blue”; multiple-category prediction = “green”; feature conjunction prediction = “brown.”

The low-category-salience condition followed the same procedure, except in experimental instructions and appearance of the categories. The category exemplars had exactly the same statistical structure and appearance as the “inner structures” in the viral stimuli in the high-salience condition but were presented without the viral outline. Participants were instructed to learn about two kinds of computer-generated drawings, with the respective categories given neutral labels (Set A, Set B). The test instance for each problem was described as a new drawing generated by the computer that could belong to either of the presented sets. In other respects, the procedure and scoring of responses were identical to those in Experiment 1.

Results and discussion

Participants in both category salience conditions always identified the target category correctly. Both groups gave modest confidence ratings for these judgments (high salience, $M = 61.40$, $SD = 3.70$; low salience, $M = 69.90$, $SD = 15.85$). Confidence did not differ as a function of category salience, $t(22) = 1.81$, $p = .08$. Hence, participants in both salience conditions understood that the category membership of the target instance was uncertain.

The proportion of feature predictions consistent with feature conjunction reasoning was at or close to ceiling in the low-salience ($M = 1.0$) and high-salience ($M = .98$, $SD = .07$) conditions. A small number of single-category predictions were made in the high-salience condition ($M = .02$, $SD = .07$). No predictions based on multiple-category reasoning were found. Across conditions, the proportion of predictions based on feature conjunction reasoning was above chance, $t(23) = 63.32$, $p < .001$, and the proportion of predictions based on single-category reasoning was again close to floor. Confidence in predictions based on feature conjunction reasoning was high and did not differ across salience conditions (high salience, $M = 84.13$, $SD = 19.11$; low salience, $M = 93.02$, $SD = 12.40$), $t(22) = 1.35$, $p = .19$. When analyzed at the individual level,

all the participants were found to make consistent use of feature conjunction reasoning.

Overall, these results show that even when steps were taken to highlight the significance of category-level information, feature conjunction remained the dominant strategy for making feature predictions under conditions of uncertain category membership.

Experiment 4

The present experiment again compared feature conjunction reasoning with single- and multiple-category approaches but provided a more precise comparison between single- and multiple-category reasoning. In Experiment 1, the test of multiple-category reasoning was relatively insensitive. The probability that the test exemplar contained the feature value most likely to be predicted by the multiple-category strategy was only marginally larger than the probability of the next most likely feature value, meaning that it was hard to distinguish between predictions based on the two approaches (see Table 1). Note that this does not alter the main conclusions from that experiment. The feature value most often chosen in Experiment 1 was predicted by feature conjunction reasoning, and this feature was assigned the lowest probability by both category-based strategies. Nonetheless, in Experiment 4, we reexamined the three reasoning strategies, using a statistical structure that ensured a clearer separation between the feature predictions of the multiple- and single-category approaches with at least a .1 difference in the predicted probabilities of respective features (see Table 3 for worked examples).

The aim of this experiment was also to extend the results of the earlier experiments in two ways. First, we included induction problems where three categories (one target, two nontarget) were relevant to feature predictions. These more complex problems were more similar to those employed in previous studies of uncertain induction (e.g., Murphy & Ross, 1994). Increasing the number of categories also increases the number of exemplars that need to be consulted to generate a feature conjunction prediction. Therefore, the

Table 3 Experiment 4: Example of probabilities of feature choices on induction test based on each induction strategy. The most likely feature choice for each strategy is given in bold (Top panel = Two-

category problems in Fig. 4 and a test exemplar with given feature “trapezoid”; Bottom panel = Three category problems in Fig. 4 and a test exemplar with feature “star”).

Feature alternatives	Single category (Murphy & Ross, 1994)	Multiple category (Anderson, 1991)	Feature conjunction
Two category problems			
$p(\text{respond blue})$	$= p(\text{blue} \text{Cleo})$ $= 4/10$ $= \mathbf{0.4}$	$= [p(\text{Cleo} \text{trapezoid}) * p(\text{blue} \text{Cleo})] +$ $[p(\text{Victoria} \text{trapezoid}) * p(\text{blue} \text{Victoria})]$ $= [0.6 * 0.4] + [0.4 * 0.1]$ $= 0.28$	$= p(\text{blue} \text{trapezoid})$ $= 0.2$
$p(\text{respond yellow})$	$= p(\text{yellow} \text{Cleo})$ $= 3/10$ $= 0.3$	$= [p(\text{Cleo} \text{trapezoid}) * p(\text{yellow} \text{Cleo})] +$ $[p(\text{Victoria} \text{trapezoid}) * p(\text{yellow} \text{Victoria})]$ $= [0.6 * 0.3] + [0.4 * 0.5]$ $= \mathbf{0.38}$	$= p(\text{yellow} \text{trapezoid})$ $= 0.3$
$p(\text{respond green})$	$= p(\text{green} \text{Cleo})$ $= 2/10$ $= 0.2$	$= [p(\text{Cleo} \text{trapezoid}) * p(\text{green} \text{Cleo})] +$ $[p(\text{Victoria} \text{trapezoid}) * p(\text{green} \text{Victoria})]$ $= [0.6 * 0.2] + [0.4 * 0.3]$ $= 0.24$	$= p(\text{green} \text{trapezoid})$ $= \mathbf{0.4}$
Three category problems			
$p(\text{respond brown})$	$= p(\text{brown} \text{John})$ $= 4/10$ $= \mathbf{0.4}$	$= [p(\text{John} \text{star}) * p(\text{brown} \text{John})] + [p(\text{Sam} \text{star})$ $* p(\text{brown} \text{Sam})] + [p(\text{Brett} \text{star}) * p(\text{brown} $ $\text{Brett})]$ $= [0.46 * 0.4] + [0.31 * 0.1] + [0.23 * 0.1]$ $= 0.24$	$= p(\text{brown} \text{star})$ $= 0.15$
$p(\text{respond aqua})$	$= p(\text{aqua} \text{John})$ $= 2/10$ $= 0.2$	$= [p(\text{John} \text{star}) * p(\text{aqua} \text{John})] + [p(\text{Sam} \text{star}) *$ $p(\text{aqua} \text{Sam})] + [p(\text{Brett} \text{star}) * p(\text{aqua} \text{Brett})]$ $= [0.46 * 0.2] + [0.31 * 0.5] + [0.23 * 0.7]$ $= \mathbf{0.41}$	$= p(\text{aqua} \text{star})$ $= 0.31$
$p(\text{respond purple})$	$= p(\text{purple} \text{John})$ $= 2/10$ $= 0.2$	$= [p(\text{John} \text{star}) * p(\text{purple} \text{John})] + [p(\text{Sam} \text{star})$ $* p(\text{purple} \text{Sam})] + [p(\text{Brett} \text{star}) * p(\text{purple} $ $\text{Brett})]$ $= [0.46 * 0.2] + [0.31 * 0.3] + [0.23 * 0.1]$ $= 0.21$	$= p(\text{purple} \text{star})$ $= \mathbf{0.46}$

dominance of feature conjunction may decrease with additional relevant categories. Alternately, the complexity of computations involved in predictions based on the multiple-category approach will increase with the number of categories that have to be considered. This may lead to a shift toward predictions based on the target category alone (i.e., single-category reasoning). A second extension involved the use of categories where features were instantiated as text descriptions (of the symptoms of fictional diseases), rather than as perceptual features.

Method

Participants

Fifty-six undergraduates participated for course credit in an introductory psychology course ($M = 20.18$ years). Participants were randomly allocated to the color/shape ($n = 27$) or text feature ($n = 29$) condition.

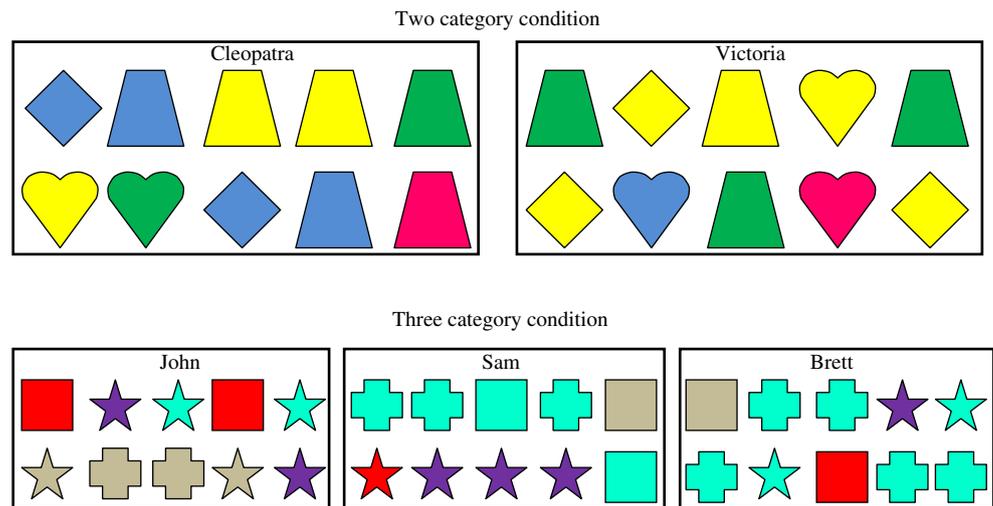
Materials

The induction problems in this experiment resembled those in the earlier experiments. Induction problems involved either two categories (as in the previous experiments) or three categories. Examples of each type are given in Fig. 4. Another modification was that the features of category exemplars were presented as color/shape combinations (as in the previous “children’s drawings” stimuli) or as text descriptions of the symptoms of fictitious diseases (see the Appendix, Fig. A2, for an example).

Procedure

Participants were assigned to either the color/shape or the text feature condition. In each feature condition, participants completed four problems that involved two categories and four problems that involved three categories. Presentation of these two problem types was blocked, with the order of block presentation randomized across participants.

Fig. 4 Example of categories from the two-category and three-category conditions in Experiment 4. For the two-category condition, the critical question was “I have a drawing of a trapezoid. What color do you think it is?” Single-category prediction = “blue”; multiple-category prediction = “yellow”; feature conjunction prediction = “green.” For the three-category condition, the critical question was “I have a drawing of a star. What color do you think it is?” Single-category prediction = “brown”; multiple-category prediction = “aqua”; feature conjunction prediction = “purple.”



The procedure for administering problems was similar to that in the earlier experiments, except that categories were presented on a 21-in. computer monitor using Runtime Revolution 3.0. The relative position of the categories on the screen was randomized across problems. As in the earlier experiments, the problems were structured so that, for the feature prediction question, the three induction strategies favored a different response alternative. However, in this experiment, there was a larger separation of the relative probabilities associated with the single- and multiple-category strategies than in Experiment 1 or 3 (see Table 3).

Results and discussion

As in the previous experiments, the accuracy of target category identification was high in all conditions ($M = .98$ correct). It did not vary as a function of the number of categories presented or surface appearance (geometric vs. text) of features ($F_s < 1$). Once again, confidence in category identification was modest ($M = 61.48$, $SD = 11.81$). Confidence ratings were unaffected by feature appearance ($F < 1.5$) but were higher when there were only two ($M = 63.52$), rather than three ($M = 59.56$), categories to be considered, $F(1, 54) = 7.78$, $p < .01$.

The proportion of feature predictions consistent with each reasoning strategy was calculated for each participant for induction problems for which the target category was correctly identified. Once again, feature conjunction was the dominant strategy across problem conditions ($M = .85$, $SD = .21$), with an overall rate of responding that was well above a chance value of .25, $t(26) = 14.50$, $p < .001$. The rates of multiple-category ($M = .14$, $SD = .17$) and single-category ($M = .07$, $SD = .13$)

reasoning were slightly higher than those in the previous experiments, but in both cases, the proportions were reliably lower than chance ($p_s < .001$).

One-way analyses of variance confirmed that feature appearance did not affect the prevalence of any of the reasoning strategies ($F_s < 1.5$). The number of categories presented had no effect on the prevalence of feature conjunction or single-category reasoning ($F_s < 1.5$). Notably, the rates of multiple-category reasoning were reliably higher for two-category problems ($M = .16$) than for three-category problems ($M = .11$), $F(1, 54) = 4.58$, $p = .04$. Confidence in predictions was moderate ($M = 57.59$, $SD = 15.65$) and did not differ across induction strategy or experimental condition, $F_s < 2.5$.

Consistent use of a reasoning strategy by individual participants was defined in the same way as in the earlier experiments. In the color/shape condition, 24 participants consistently used feature conjunction reasoning, 1 used multiple-category reasoning, and 2 participants did not use any strategy consistently. In the text condition, 22 participants consistently used feature conjunction, and 7 did not use any consistent strategy.

This experiment provided a more valid simultaneous comparison of feature conjunction with each of the category-based approaches to feature prediction with uncertain categories. However, this had little effect on the main finding; predictions were usually based on feature conjunction and rarely on category membership (either single or multiple). Increasing the number of relevant categories did not affect the dominance of feature conjunction, although there was some evidence that it shifted people in the direction of single- as opposed to multiple-category reasoning. This experiment also showed that feature conjunction is equally prevalent in categories instantiated with text-based, as well as perceptual, features.

General discussion

In the present experiments, we examined how people make feature predictions about target instances whose category membership is uncertain. Previous work on this issue has focused on how people make predictions using information from multiple categories (e.g., Anderson, 1991) or just a single target category (e.g., Murphy & Ross, 1994, 2005). So far, the weight of evidence suggests that the latter approach is the one most often used across a range of stimulus and task contexts (Murphy & Ross, 2007).

Our experiments, however, challenge this conclusion by showing that a noncategorical feature conjunction approach is frequently employed when people make feature predictions with uncertain categories. [Experiment 1](#) showed that when people made feature predictions with uncertain categories, feature conjunction reasoning was overwhelmingly endorsed, with very little use of single- or multiple-category reasoning. [Experiment 2](#) showed that the use of feature conjunction involved assuming that feature dimensions were correlated *and* that this assumption applied to all available exemplars that shared a given feature with a test instance. In [Experiment 3](#), despite the fact that a number of measures were taken to increase the perceived salience and meaningfulness of category-level information, feature conjunction reasoning was still overwhelmingly endorsed. [Experiment 4](#) showed that the dominance of feature conjunction extends to more complex induction problem structures and an entirely different means of feature presentation (i.e., exemplars presented as lists of written features).

These results have important implications for the way inductive reasoning is conceptualized. When the category membership of a target instance is known with certainty, there is little doubt that feature predictions are made on the basis of category-level information (e.g., Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Rips, 1975). However, when category membership is uncertain, our experiments suggest that people will often follow a noncategorical approach to induction. Specifically, when people are presented with a feature belonging to a novel instance, they often base their predictions on this *certain featural* knowledge, rather than on uncertain category membership.

The possibility that people use feature conjunction in induction with uncertain categories has been acknowledged in past work (e.g., Murphy & Ross, 1994, 2010a). Rather than trying to control for or eliminate this approach to reasoning, our experiments are the first to clearly differentiate between predictions based on feature conjunction, single-category, and multiple-category approaches. Our finding of an overwhelming preference for feature conjunction suggests a reinterpretation of previous work on

induction with uncertain categories. In cases in which feature conjunction predictions are confounded with those of category-based strategies (e.g., Murphy & Ross, 1994, [Experiments 1-3, 2005](#)), the findings pointing to single-category reasoning may actually reflect use of feature conjunction. Other studies (e.g., Murphy & Ross, [Experiments 5 and 6](#)) that have made it difficult to employ feature conjunction reasoning have shown evidence of single-category reasoning. Our findings, however, suggest that such strategies may be prevalent only when feature conjunction reasoning is not possible.

The finding that people regularly use feature conjunction for making feature predictions with uncertain categories is consistent with models that propose that categorization proceeds via an assessment of the similarity of a novel target to known category exemplars (e.g., Medin & Schaffer, 1978; Nosofsky, 1986). Furthermore, these findings are in line with previous work showing that categorization judgments are influenced by the specific similarity between a target instance and previously encountered instances, even when people are provided with a perfectly predictive classification rule (e.g., Allen & Brooks, 1991; Regehr & Brooks, 1993). Our experiments, however, suggest that the effects of exemplar similarity extend to the domain of inductive inference.

The present experiments used artificial categories, but we see the findings as relevant to inductive prediction in the real world. Imagine, for example, that you see a new model car that has some features that are consistent with the category of sports utility vehicle (e.g., large tires, an elevated chassis) and some features that are consistent with the category of luxury vehicle (e.g., leather seats, expensive sound system). You want to predict whether the car will be suitable for driving on rugged mountain tracks. If you could confidently assign the object to one of the vehicle categories, you could make the prediction on the basis of typical category features. However, given the category uncertainty, you may resort to a prediction based on the presence or absence of features that have previously been predictive of success in mountain driving (e.g., a large-capacity engine).

This example also highlights what may be an important boundary condition for feature conjunction in the everyday world; namely, this approach assumes that a sample of relevant exemplars is readily available when feature predictions are required. In the present experiments, all relevant category information was visible throughout the prediction process. In most cases of everyday prediction, however, people will have to rely on memory for exemplars (and their features). Over long periods of time following initial learning, or if large numbers of exemplars have to be considered, it may become difficult to generate accurate predictions on the

basis of feature conjunction, and people may show greater use of category-based approaches.

Another caveat about the generalization of the present results is that despite the manipulation of category salience in [Experiment 3](#), the present experimental categories clearly had lower levels of internal coherence than did those found in basic-level natural categories. It remains an open question as to whether category-based reasoning would be more common for feature predictions about objects that might belong to two or more categories with high levels of coherence. The overwhelming endorsement of feature conjunction reasoning found in the present experiments, however, suggests that, as long as category membership is uncertain, feature conjunction may be seen as a viable approach for deriving predictions.

Implications for models of induction

One model that has been the focus of much previous work on induction under category uncertainty is Anderson's (1991) Rational model. When multiple uncertain categories can be used for future prediction, the model predicts that people will use them. As was noted earlier, so far there has been little evidence found to support this prediction (e.g., Murphy & Ross, 1994, 2010b). The implications of the present experiments for the Rational model are more complicated. A key assumption of the model is that people treat feature dimensions as conditionally independent. The present data clearly challenge this assumption. On the other hand, the model has considerable flexibility, allowing it to create exemplar clusters that optimize feature prediction. In many cases, these feature clusters will correspond to labeled categories. But in the present experiments, knowing a feature of a novel instance (and being able to check its correlation with other features) was arguably more useful for feature prediction than was knowing about its category membership. In principle, the model could explain the present findings, but only by ignoring the given categories and assigning each exemplar to its own cluster. In effect, this would make the Rational model function like an exemplar model (see Nosofsky, 1991).

An alternative theoretical approach was suggested by Ross and Murphy (1996), who proposed two discrete stages in the derivation of feature predictions under category uncertainty. When a test instance with a given feature is presented, the first stage involves identifying the target category. The second stage involves making a feature prediction on the basis of the output from the first stage. Although people appreciate the uncertainty of category membership during the categorization phase, only information from the target category is held in working memory during the subsequent induction phase. If alternative categories are made salient during induction, people may

use multiple-category reasoning (e.g., Hayes & Newell, 2009; Ross & Murphy, 1996), but such cases are relatively rare.

The findings of the present experiments are problematic for this model. Like the Rational model, most instantiations of single-category reasoning have involved assumptions of feature independence, so in this regard, our feature conjunction findings challenge this approach.² Moreover, [Experiment 3](#) showed that when feature conjunction was a viable method for generating predictions, there was little evidence that people focused on information from the target category, despite the fact that participants could identify the correct target category for each test item. Contrary to the Ross and Murphy (1996) model, people appeared to be thinking categorically when classifying novel instances, but not when they made feature predictions.

Conclusions

Our findings suggest that, when dealing with objects whose category membership is uncertain, people do not always use categories to make feature predictions. Instead, they often base predictions on a comparison of the features of familiar and novel exemplars across category boundaries. This is an important finding because most previous work has assumed that feature predictions about category members rely primarily on category-level information. The emphasis on category-based predictions has been driven, in part, by the use of experimental paradigms in which people are provided with the possible category labels only for a novel instance, together with some indication of their respective probabilities (e.g., Hayes & Chen, 2008; Ross & Murphy, 1996). In such cases, the question of how people make feature predictions reduces to a question of single-category or multiple-category reasoning. Our data show, however, that when the features of individual category members are available for inspection, prediction based on feature conjunction is a viable alternative to category-based prediction. It remains for future work to establish the boundary conditions of feature conjunction reasoning. In particular, this work should focus on whether feature conjunction reasoning persists when exemplars have to be retrieved from memory at the prediction stage and when categories with a more coherent internal structure are presented.

² To be fair, Murphy and Ross (1994) acknowledged that the assumption of feature independence is not critical to their single-category approach. Nevertheless, [Experiment 2](#) is the first attempt to differentiate the predictions of feature conjunction based on single or multiple categories.

Appendix

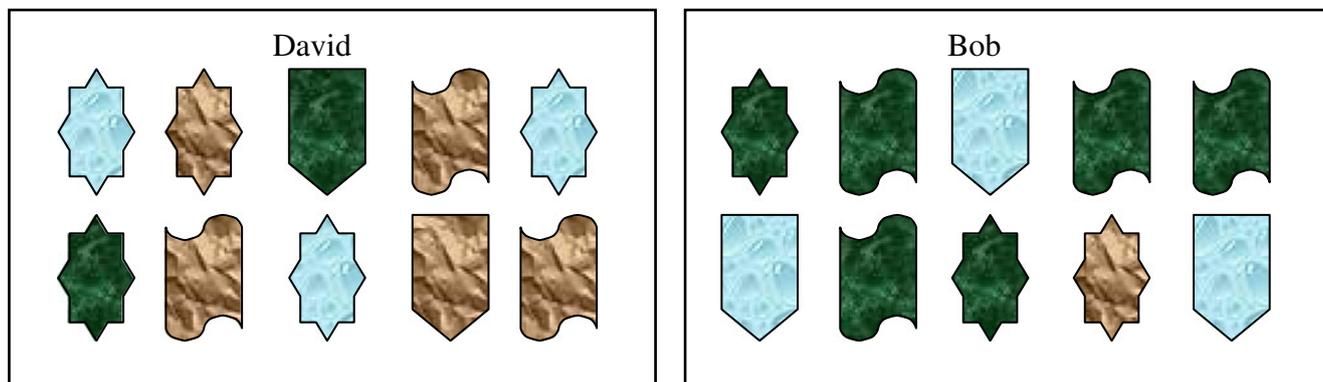


Fig. A1 An example of a pattern/shape category (i.e. “graphic design student drawings”) used in Experiment 1

Fig. A2 An example of the text-based categories used in Experiment 4 (two-category condition)

	Effertus		Laboro	
	Symptom 1	Symptom 2	Symptom 1	Symptom 2
Patient 1	Chest Congestion	Bloating	Sneezing	Vomiting
Patient 2	Breathlessness	Difficulty Swallowing	Chest Congestion	Vomiting
Patient 3	Chest Congestion	Stomach Ache	Breathlessness	Vomiting
Patient 4	Chest Congestion	Difficulty Swallowing	Chest Congestion	Difficulty Swallowing
Patient 5	Chest Congestion	Vomiting	Chest Congestion	Difficulty Swallowing
Patient 6	Breathlessness	Vomiting	Chest Congestion	Difficulty Swallowing
Patient 7	Chest Congestion	Bloating	Sneezing	Vomiting
Patient 8	Chest Congestion	Vomiting	Breathlessness	Bloating
Patient 9	Sneezing	Bloating	Sneezing	Vomiting
Patient 10	Sneezing	Bloating	Breathlessness	Stomach Ache

References

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

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409–429.

Chin-Parker, S., & Ross, B. H. (2002). The effect of category learning on sensitivity to within-category correlations. *Memory & Cognition*, 30, 353–362.

Hayes, B. K., & Chen, T. H. J. (2008). Clinical expertise and reasoning with uncertain categories. *Psychonomic Bulletin and Review*, 15, 1002–1007.

Hayes, B. K., & Newell, B. R. (2009). Induction with uncertain categories: When do people consider all the alternatives? *Memory & Cognition*, 37, 730–743.

Hayes, B. K., Ruthven, C., & Newell, B. R. (2007). Inferring properties when categorization is uncertain: A feature-conjunction account. In D. McNamara & G. Tafton (Eds.), *Proceedings of the 28th Meeting of the Cognitive Science Society* (pp. 209–215). Hillsdale: Erlbaum.

- Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, 7, 569–592.
- Gelman, S. A. (2003). *The essential child*. New York: Oxford University Press.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207–238.
- Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, 27, 148–193.
- Murphy, G. L., & Ross, B. H. (2005). The two faces of typicality in category based induction. *Cognition*, 95, 175–200.
- Murphy, G. L., & Ross, B. H. (2007). Use of single or multiple categories in category-based induction. In E. Heit & A. Feeney (Eds.), *Inductive reasoning: Experimental, developmental and mathematical approaches*. Cambridge University Press.
- Murphy, G. L., & Ross, B. H. (2010a). Category vs. object knowledge in category-based induction. *Journal of Memory and Language*. doi:10.1016/j.jml.2009.12.002
- Murphy, G. L., & Ross, B. H. (2010b). Uncertainty in category-based induction: When do people integrate across categories? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 263–276.
- Murphy, G. L., & Wisnewski, E. J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications* (Vol. 2, pp. 23–45). Chichester: Ellis Horwood.
- Newell, B. R., Paton, H., Hayes, B. K., & Griffiths, O. (2010). *Speeded induction under uncertainty: The influence of multiple categories and feature conjunctions*. Manuscript under review.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nosofsky, R. M. (1991). Relation between the rational model and the context model of categorization. *Psychological Science*, 2, 416–421.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, 97, 185–200.
- Patalano, A. L., Chin-Parker, S., & Ross, B. H. (2006). The importance of being coherent: Category coherence, cross-classification, and reasoning. *Journal of Memory and Language*, 54, 407–424.
- 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.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behaviour*, 14, 665–681.
- Ross, B. H., & Murphy, G. L. (1996). Category-based predictions: Influence of uncertainty and feature associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 736–753.
- Thibaut, J. P., & Gelaes, S. (2006). Exemplar effects in the context of a categorization rule: Featural and holistic influence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1403–1415.
- Thomas, R. D. (1998). Learning correlations in categorization tasks using large, ill-defined categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 119–143.
- Verde, M. F., Murphy, G. L., & Ross, B. H. (2005). Influence of multiple categories on the prediction of unknown properties. *Memory & Cognition*, 33, 479–487.
- Yamauchi, T., & Markman, A. B. (1998). Category learning by inference and classification. *Journal of Memory and Language*, 39, 124–148.
- Yamauchi, T., Love, B. C., & Markman, A. B. (2002). Learning nonlinearly separable categories by inference and classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 2585–2593.

Christopher Papadopoulos, Brett K. Hayes and Ben R. Newell, School of Psychology, University of New South Wales, Sydney, Australia. This work was supported by an Australian Postgraduate Award to the first author and Australian Research Council Discovery Grant DP0770292 to the second and third authors. We would like to thank Oren Griffiths for his valuable comments on this research, as well as Carissa Bonner and Helen Paton for their assistance with programming.