Where to look first for an explanation of induction with uncertain categories

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

Published online: 21 September 2011
© Psychonomic Society, Inc. 2011

Abstract Recent research has examined how people predict unobserved features of an object when its category membership is ambiguous. The debate has focused on whether predictions are based solely on information from the most likely category, or whether information from other possible categories is also used. In the present experiment, we compared these category-based approaches with feature conjunction reasoning, where predictions are based on a comparison among exemplars (rather than categories) that share features with a target object. Reasoning strategies were assessed by examining patterns of feature prediction and by using an eye gaze measure of attention during induction. The main findings were (1) the majority of participants used feature conjunction rather than categorical strategies, (2) people predominantly gazed at the exemplars that were most similar to the target object, and (3) although people gazed most at the most probable category to which an object could belong, they also attended to other plausible category alternatives during induction. These findings question the extent to which category-based reasoning is used for induction when category membership is uncertain.

Keywords Inductive reasoning · Category-based reasoning · Uncertain categorization

A core function of categories is to make predictions about unobserved properties of category members, but relatively little is known about how people make these predictions when an object's category membership is uncertain

(Murphy & Ross, 2007). This is an important issue, because people often have to make inductive predictions when category membership is ambiguous. Imagine, for example, that you see a new model car that has many of the characteristic features of the category “sports utility vehicle” (SUV) (e.g., large tires and an elevated chassis). However, it also has features that are consistent with the category of “luxury vehicle” (e.g., leather seats and an expensive sound system), so its category membership remains ambiguous.

If you wanted to make a prediction about some unobserved properties of the car (e.g., whether it is suitable for driving on rugged mountain tracks), a number of alternative inductive strategies are possible. Some strategies involve consideration of category membership. According to the “multiple-category” approach, you would first identify the various categories to which the car might belong (e.g., SUV or luxury vehicle) and work out the conditional probability of the predicted feature in each category (see Anderson, 1991, for details). These probabilities are then weighted according to the likelihood of membership in the respective categories and summed across category alternatives. An alternative, “single-category” approach involves identifying the most likely or target category for the given object (e.g., SUV) and then making predictions *as if* the object was known to belong to this category (see, e.g., Murphy & Ross, 1994). Although the process of hedging across multiple categories is normative according to Bayesian calculus, numerous experimental studies have shown that most people prefer to make predictions based on the single-category heuristic (see Hayes, Heit, & Swendsen, 2010, and Murphy & Ross, 2007, for reviews).

Note that both category-based accounts assume that people treat the known features of a test instance and the

O. Griffiths (✉) · B. K. Hayes · B. R. Newell · C. Papadopoulos
University of New South Wales,
Sydney, Australia
e-mail: oren.griffiths@unsw.edu.au

unknown features that have to be predicted as *conditionally independent*. This means that information about a given feature is used only to assign relative weights to target and nontarget categories (multiple-category approach) or to determine which category should be the focus of attention (single-category approach). Subsequently, the given feature plays no direct role in prediction.

More recently, it has been suggested that people make more use of featural information in inductive predictions about specific objects than is suggested by either of the category-based approaches (see, e.g., Murphy & Ross, 2010; Newell, Paton, Hayes, & Griffiths, 2010; Papadopoulos, Hayes, & Newell, 2011). The general argument is that relying on knowledge about typical category features may be insufficient when we have to make predictions about individual cases, especially when those cases do not fit neatly into a single known category. In such cases, people may base their predictions on relationships between features rather than on category membership. In our ambiguous car example, rather than agonizing over whether it is really an SUV or a luxury car with SUV trappings, one might be better off making predictions about its driving capabilities based on knowledge about specific features (e.g., engine capacity). In such cases, rather than treating features as conditionally independent, people use their knowledge of feature correlations as the basis for making inductive predictions.

Evidence supporting this feature-based approach was reported by Papadopoulos, Hayes, and Newell (2011) using categories similar to those in Fig. 1. As in previous studies (e.g., Murphy & Ross, 1994), each category was presented as a sample of drawings done by different children. Participants were then shown part of a new drawing with a given feature (e.g., a triangle) and asked to predict an unobserved feature (color). Note that the category membership of this test exemplar is uncertain. Although one child (Barbara) drew the most triangles and was the target

category, others (Patricia and Mary) drew some triangles and could also be considered. Critically, the stimuli were constructed so that the category-based and feature-based approaches yielded different predictions. According to Bayesian multiple-category reasoning (Anderson, 1991), the most likely color, based on a weighted additive sum of the typical colors in the three category alternatives (Barbara, Patricia, and Mary), is “red.” Consideration of the typical features within the target category of Barbara alone (the single-category approach) would lead to the prediction of “yellow.” Alternately, if people eschewed categories and only examined co-occurrences between the given feature (triangle) and various colors across all of the available exemplars, then the prediction would be “green” (see Table 1 for more detailed derivations of the predictions of each strategy).

Papadopoulos et al. (2011) found that such “feature conjunction” predictions were preferred over those derived from category-based induction by a large majority of participants (see also Murphy & Ross, 2010). Indeed, they found that almost all participants reliably used a conjunction strategy across several different inductive problems. Furthermore, conjunction strategies were favored even when steps were taken to highlight the internal coherence of the experimental categories (e.g., by increasing the within-category similarity of exemplars). Following these observations, we anticipated that most participants would select a particular conjunction strategy and consistently use it to make feature predictions.

The present study aimed to further examine the use of feature conjunction, as opposed to category-based, approaches for predictions about objects with uncertain category membership. One issue that remains contentious is the range of exemplars consulted when people make feature-based predictions. Papadopoulos et al. (2011) found that people made feature conjunction predictions based on

Fig. 1 A typical stimulus set. Participants were asked to make an inference about an unobserved property (in this case, color) of a given object (e.g., a triangle). There is a clear target category (Barbara), two other relevant categories (Patricia and Mary), and an irrelevant control category (Linda)

Table 1 Feature predictions example calculations of probabilities of feature predictions using each of the four induction strategies

Responses	Single Category	Multiple Category	Single-Category Conjunction	Multiple-Category Conjunction
$p(\text{yellow})$	$= p(\text{yellow} \text{Barbara})$ $= 5/10$ $= .5$	$= p(\text{yellow} \text{Barbara}) * p(\text{Barbara} \Delta)$ $+ p(\text{yellow} \text{Patricia}) * p(\text{Patricia} \Delta)$ $+ p(\text{yellow} \text{Mary}) * p(\text{Mary} \Delta)$ $= (.5 * .5) + (0 * .25) + (0 * .25)$ $= .25$	$= p(\text{yellow} \text{Barbara} \& \Delta)$ $= 2/6$ $= .33$	$= p(\text{yellow} \Delta)$ $= 2/12$ $= .17$
$p(\text{red})$	$= p(\text{red} \text{Barbara})$ $= 1/10$ $= .1$	$= p(\text{red} \text{Barbara}) * p(\text{Barbara} \Delta)$ $+ p(\text{red} \text{Patricia}) * p(\text{Patricia} \Delta)$ $+ p(\text{red} \text{Mary}) * p(\text{Mary} \Delta)$ $= (.1 * .5) + (.7 * .25) + (.7 * .25)$ $= .4$	$= p(\text{red} \text{Barbara} \& \Delta)$ $= 0/6$ $= 0$	$= p(\text{red} \Delta)$ $= 0/12$ $= 0$
$p(\text{blue})$	$= p(\text{blue} \text{Barbara})$ $= 4/10$ $= .4$	$= p(\text{blue} \text{Barbara}) * p(\text{Barbara} \Delta)$ $+ p(\text{blue} \text{Patricia}) * p(\text{Patricia} \Delta)$ $+ p(\text{blue} \text{Mary}) * p(\text{Mary} \Delta)$ $= (.4 * .5) + (0 * .25) + (0 * .25)$ $= .2$	$= p(\text{blue} \text{Barbara} \& \Delta)$ $= 4/6$ $= .67$	$= p(\text{blue} \Delta)$ $= 4/12$ $= .33$
$p(\text{green})$	$= p(\text{green} \text{Barbara})$ $= 0/10$ $= 0$	$= p(\text{green} \text{Barbara}) * p(\text{Barbara} \Delta)$ $+ p(\text{green} \text{Patricia}) * p(\text{Patricia} \Delta)$ $+ p(\text{green} \text{Mary}) * p(\text{Mary} \Delta)$ $= (0 * .5) + (.3 * .25) + (.3 * .25)$ $= .15$	$= p(\text{green} \text{Barbara} \& \Delta)$ $= 0/6$ $= 0$	$= p(\text{green} \Delta)$ $= 6/12$ $= .5$

This example uses the category structure shown in Fig. 1. Participants are told that they have a triangle and are asked to predict the most likely color of that triangle. The most likely feature prediction for each strategy is shown in bold

an examination of *all* visible exemplars, including those from both target and nontarget categories (“multiple-category feature conjunction”). In contrast, Murphy and Ross (2010) focused on feature conjunctions based on correlations between features *within the target category* (“single-category feature conjunction”). In Fig. 1, for example, if the given feature was triangle, the single-category version of feature conjunction would focus on the triangles drawn by Barbara (leading to the feature prediction “blue”).

The present experiment therefore sought to extend the study of induction with uncertain categories in two ways. First, as well as comparing category-based with feature-based reasoning, this study compared reasoning based on the single- and multiple-category variants of feature conjunction. This was accomplished by developing induction problems like that shown in Fig. 1, where the four possible reasoning strategies (multiple-category, single-category, multiple-category conjunction, single-category conjunction) led to different feature predictions. Previous studies have only compared a subset of these four possible approaches. For example, although Murphy and Ross (2010) favored the single-category version of feature conjunction, their paradigm did not allow for a direct comparison of single- and multiple-category conjunction strategies. Papadopoulos et al. (2011, Exp. 2) did compare these two approaches (and found that multiple-category conjunction was preferred) but did not place the two feature conjunction approaches in competition with category-based reasoning. Hence, this study is the first to simultaneously compare the use of all four strategies for reasoning with uncertain categories; that is, both single- and multiple-category versions of category-based and feature-based approaches.

The second innovation in this study was the use of eye gaze measures to examine how attention is distributed during induction with uncertain categories. This allowed us to test contrasting accounts of the relations between attention and decision processes in induction. Past discussion of induction strategies has implied that the differential weighting of exemplars in the decision process is a product of previous patterns of attentional allocation to different categories and exemplars. For example, Ross and Murphy (1996) suggested that after a test stimulus is assigned to a target category, only exemplars within that category are attended to during the induction process. Assuming that eye gaze is systematically related to attention (cf. Deubel & Schneider, 1996), each of the four reasoning strategies outlined above would predict a different pattern of attentional allocation (see Fig. 2 for a summary and worked example). For example, the single-category account suggests that participants first identify the target category and then focus attention primarily on target category members when making a feature prediction. In

contrast, the multiple-category account predicts that participants will also gaze at exemplars from the nontarget categories throughout the induction procedure, perhaps in proportion to the likelihood that the test object belongs to each category. People may gaze longer at target than nontarget categories, but importantly, the multiple-category approach predicts that participants will look longer at the nontarget exemplars than at the irrelevant control category (see row 2 in Fig. 2).

Both feature conjunction approaches predict that, during induction, people will spend more time gazing at exemplars with the given feature (e.g., triangles) than on exemplars with a different value on the given dimension (e.g., squares). However, according to the multiple-category conjunction account, this pattern should be found in both target and nontarget categories. In contrast, single-category conjunction predicts that attention will be restricted to exemplars with the given feature that belong to the target category (see the bottom rows of Fig. 2).

However, it is also possible that these induction strategies are not implemented in attentional behavior, and instead reflect differential weighting within the decision process once all exemplars have been considered. That is, all exemplars may be examined during the induction task, but then a privileged few form the basis for subsequent inductive prediction. If this were the case, there should be little correspondence between patterns of eye gaze and feature prediction.

The eye gaze measures therefore go beyond the feature prediction data by indicating the locus of the selective attention process underlying each inductive strategy. Furthermore, if people consistently use a particular inductive strategy, as seen in Papadopoulos et al.’s (2011) study, it may also be possible to examine whether such strategies exhibit characteristic attentional (eye gaze) profiles.

Method

Participants

A group of 48 introductory psychology students participated for course credit. Their mean age was 18.98 years, and 34 participants were females.

Design and materials

The eight induction problems comprised four categories, each containing ten exemplars. Each exemplar could vary on three dimensions (color, shape, and pattern), and each dimension had four levels (as shown in Fig. 1). The category structures were constructed so that, for a given test

Fig. 2 Types of stimuli used to perform feature induction under each of the four induction strategies, and the corresponding predictions for eye gaze behavior generated from these strategies. Stimuli are sorted according to their categories (target, nontarget, control) and whether they possess the given feature (e.g., are triangular, Δ) or not ($\sim\Delta$).

Items in bold are important for that particular strategy. The height of each column represents the predicted relative amount of time spent gazing at each type of stimulus, where high columns indicate longer gaze duration (in arbitrary units)

instance, each of the four induction strategies clearly favored a different feature prediction (see Table 1). Additionally, each strategy also led to a clear “second most likely” prediction, and we sought to equate the differences in choice probabilities between the most likely and the second most likely predictions so as to minimize any differences in the ease with which each strategy could be applied.

Procedure

Participants’ heads were placed on a fixed chinrest, and a five-point gaze calibration was performed. Eight induction problems were then presented sequentially. For each problem, participants were told that they would be shown a representative sample of pictures drawn by four girls. They would then be shown a novel picture with some of the details missing. The task was to predict the missing features

of the drawing (feature prediction) and the girl who drew it (target category identification).

As is shown in Fig. 3, each problem trial began with a 30-s familiarization phase in which participants were shown the four stimulus categories. The upper part of the screen was divided into four quadrants, one per category, and the 10 exemplars were evenly spaced within that quadrant. Each exemplar measured 100 pixels square. Each category was labeled with a female name shown below the exemplars. A new, randomly chosen female name was selected for each category at the beginning of each trial. All information, instructions, and response options were shown in the lower third of the screen. The stimuli were locked onscreen for 30 s, after which participants could progress to the first test question by clicking on an onscreen button.

The four categories of geometric stimuli were then removed, and a statement describing the given feature (e.g., “I have a picture of a triangle”) was shown for 2 s. The first

test question was then shown for 2 s (either a target category identification, “Which girl do you think drew this picture?” or a feature prediction question, e.g., “What color do you think this picture is most likely to be?”). The categories then reappeared at the top of the screen, and participants indicated their classification decision by clicking on the appropriate button. They then rated their confidence in this judgment on a 1–100 scale (1 = *not at all certain*, 100 = *very certain*). The same procedure (shown in Fig. 3) was then repeated for the second question. To control for the effect of question order on inductive feature predictions (see, e.g., Hayes & Newell, 2009), the order of the classification and induction questions was reversed for half of the participants.

The assignment of all stimulus variables (e.g., color, shape, patterning) as well as of the screen positions of categories (target, nontargets, control) was randomly determined on each trial.

Eye gaze measurement

Eye gaze was recorded using a Tobii T60 eyetracker, with the fixation filter set to a radius of 20 pixels and a minimum dwell time of 40 ms. Gaze recording commenced when the induction question and response options were shown (see Fig. 3) and terminated once a feature prediction was made. The screen location, duration, and time of occurrence of each fixation during this period were recorded. The total time spent fixating on each exemplar in the prereponse period was recorded, as was the total time spent fixating on any part of the screen in that period. From these fixation tallies, we then calculated the time spent fixating on each

exemplar, expressed as a proportion of the total time spent fixating on any part of the screen (hereafter, “dwell time”; for further calculation details, see the [Appendix](#)). These fixation proportion scores were then averaged amongst exemplars of the same type, where the five types of exemplars were as follows: 6 exemplars in the target category with the given feature (labeled “TwitF”), 4 exemplars in the target category without the given feature (“TnoF”), 6 exemplars in the nontarget categories with the given feature (“NTwithF”), 14 exemplars in the nontarget categories without the given feature (“NTnoF”), and 10 exemplars in the control category, none of which possessed the given feature (“control”).

Results

Identification of the target category was generally accurate (90% correct). The mean confidence for target identification was 65.24 (out of 100). This modest figure suggests that participants recognized the uncertainty about the category memberships of the test items.

Both the category-based approaches to induction and single-category conjunction are predicated on accurate identification of the target. Hence, the trials on which the target category was incorrectly identified were removed from further analysis. One participant was removed because he failed to identify the target category on the majority of trials.

A different prediction option on each trial was favored by each of the four induction strategies. The responses consistent with each strategy were tallied across trials to produce a score for each strategy (out of 8). Figure 4 shows

Fig. 3 Summary of the experimental procedure. Eye gaze was recorded during the second-to-last step (i.e., from the presentation of the induction question until an induction response was made). The panels superimposed upon the directional arrow offer a magnified view of the text shown on each screen

Fig. 4 Mean number of responses consistent with each of the four induction strategies. The broken line indicates the chance (uniform) pattern of judgment. The columns in Fig. 4 sum to less than 8 due to the exclusion criteria

that feature predictions based on conjunction strategies were preferred over category-based predictions.

Patterns of individual responding were examined by classifying participants according to their dominant feature prediction choices. Those who predicted the feature associated with a particular strategy on 5 or more trials (out of 8) were considered “consistent” users of that strategy. By chance alone, approximately 10.92% of participants ought to be classified as a consistent user of a strategy. Under this criterion, 10 participants consistently used single-category conjunction, 23 used multiple-category conjunction, and 14 did not show a consistent pattern. No consistent users of either single- or multiple-category reasoning were observed. Although the present rates of consistent strategy use were not as high as those observed previously (Papadopoulos et al., 2011), a binomial test revealed that the frequencies of classification as either a single-conjunction user or a multiple-conjunction user were both above what would be expected by chance (single category, $p = .03$; multiple category, $p < .001$). The rates of strategy use did not appear to depend on any of the stimulus variables (e.g., dimension of the given feature), although there appeared to be a relationship between strategy choice and question order consistent with that shown by Hayes and Newell (2009). Namely, 9 of the 10 single-category conjunction users received the category identification question before the feature induction question (as compared to the 10 out of 23 multiple-category conjunction users who received

this question order). However, cell sizes were too small for a valid chi-square analysis of these question order effects.

Strategy use and eye gaze

The gaze behavior of the two consistent user groups (single- and multiple-category conjunction users) was analyzed using a series of planned orthogonal contrasts in a multivariate ANOVA (see Fig. 5). This analysis excluded the 14 inconsistent responders.¹ The type of exemplar (TwthF, TnoF, NTwthF, TnoF, or control) was a within-subjects factor, and group assignment (multiple-conjunction or single-conjunction user) was a between-subjects factor. Note that one consequence of dividing the individual exemplars into the five types of exemplars was that group main effects were uninterpretable. The influence of group assignment could only be determined by examining when they interact with the contrasts that compare dwell times between specific types of exemplars. Hence, only main effects based on planned comparisons between the various exemplar types and the group $\times$ exemplar type interactions were examined.

As anticipated, participants gazed longer at exemplars from the target category (TwthF and TnoF) than at exemplars in the other categories (NTwthF, NTnoF, and control), $F(1, 31) = 158.14$, $p < .001$, $\eta^2 = .53$. This is consistent with all induction strategies. The bias toward the target category exemplars was stronger for the single-category conjunction group than for the multiple-category conjunction group, $F(1, 31) = 28.28$, $p < .001$, $\eta^2 = .1$. Within the target category, participants gazed at the exemplars with the given feature (TwthF) more than at those without that feature (TnoF), $F(1, 31) = 25.89$, $p < .001$, $\eta^2 = .07$. Amongst the other categories, exemplars in the two plausible alternative categories (NTwthF and NTnoF) were gazed at longer than those in the control category, $F(1, 31) = 74.90$, $p < .001$, $\eta^2 = .03$. This shows that people spent a nontrivial amount of time gazing at the nontarget categories, as well as at the target category. A trend was observed in which this bias for the nontarget over the control exemplars was larger for multiple-category conjunction users than for single-category conjunction users, $F(1, 31) = 4.02$, $p = .05$, $\eta^2 < .01$. Within the nontarget categories, exemplars with the given feature (NTwthF) were gazed at longer than those without that feature (NTnoF), $F(1, 31) = 15.12$, $p < .001$, $\eta^2 = .02$. No further exemplar type effects or interactions reached significance.

¹ A corresponding analysis including all participants was also carried out. This analysis also found a gaze preference for exemplars in the target category and for exemplars with the given feature.

Fig. 5 Mean proportion of gaze time per exemplar, as a function of the use of single-category conjunction or multiple-category conjunction reasoning. The x-axis refers to the type of exemplar: “T” = a member of the target category, “NT” = a member of the non-target category, “withF” = an exemplar with the given feature, and “noF” = an exemplar without that feature

We also examined whether eye gaze allocation changed during the course of prediction by dividing dwell times in each trial into three equal time bins (early-, mid-, and late-trial). A repeated measures analysis of dwell times across trial times found that the gaze biases reported above were generally stable throughout prediction, although the bias toward the target category diminished across the trial, $F(1, 31) = 5.95$, $p = .02$, $\eta^2 < .01$. This effect did not vary across the two groups of conjunction reasoners.

Discussion

When faced with making feature predictions about objects whose category membership is uncertain, we found that people often based these predictions on relations between features rather than on category membership. These results are consistent with other recent studies that have shown a preference for feature-based over category-based reasoning (Murphy & Ross, 2010; Newell et al., 2010; Papadopoulos et al., 2011). That is, people generated feature predictions based on the frequency with which the given feature co-occurred with other features.

Insights from eyetracking

The eye gaze data shed light on the attentional mechanisms involved in feature prediction. Contrary to strong accounts of the single-category and single-category conjunction approaches (e.g., Ross & Murphy, 1996), people did not focus their search exclusively on the exemplars in the target category after it had been identified. In fact, the time bin analysis showed that attention to the target category decreased across the prediction process. This suggests that

most participants (even single-category conjunction users) considered exemplars from both target and nontarget categories prior to forming an inductive judgment, rather than just focusing their search behaviors on the most likely category. The eye gaze data also confirmed that prior to making a conjunction-based prediction, people selectively attended to exemplars that possessed the given feature and to exemplars in the target category. This observation suggests that at least some of the weighting process, in which exemplars with the given feature and those in the target category are preferentially processed, is evident in participants' attentional behavior. That is, exemplar weighting is not confined to the decision stage in the induction process.

Given that some exemplars are preferentially attended to during training, one may wonder about the extent to which this preference reflects the attentionally weighted similarity comparison process inherent in some exemplars models of categorization (e.g., Nosofsky, 1986). Interestingly, the eye gaze data are broadly consistent with the attentional parameters required to produce conjunction-based inductive judgments. For example, a strong bias toward exemplars with the given feature (high weighting on the given dimension) would lead to a conjunction-based inductive judgment, as observed. Furthermore, the size of the attentional bias toward the target category (the weight assigned to the category label dimension) ought to determine whether inductive prediction is consistent with the single- or multiple-category conjunction strategy: those with a stronger bias ought to rely on single-category conjunction reasoning. This pattern was also seen in the present data. However, it is difficult to interpret this finding, since the direction of the relationship between feature weighting in induction and overt gaze behaviors is

ambiguous. For example, attentional weight (in exemplar models) might be instantiated as higher salience for particular features, leading participants to preferentially attend to some exemplars (e.g., those with the given feature), and this limited sampling could shape their inductive judgments. Alternatively, participants might first determine which exemplars are likely to be most important for the inductive task (via consultation of an internal representation), and this belief could then guide their search toward those exemplars. Indeed, the ability of eyetracking technology to examine hypotheses such as these remains one of its most exciting possibilities.

Dominance of multiple- over single-category feature conjunction

The present study also advanced understanding of feature-based reasoning by carrying out the first simultaneous comparison of single- and multiple-category versions of feature conjunction along with category-based induction strategies. The prediction data suggest that both kinds of conjunction-based reasoning were used, although a majority used multiple-category conjunction. However, strategy consistency was not perfect. Although most people used a particular strategy on the majority of inductive problems, almost all participants tried a second strategy on at least one problem, and some (14) did not consistently use any strategy. Moreover, the strategy use appeared to depend on question order. More people consistently used single-category conjunction when they were given the classification question before the critical induction question. As argued by Hayes and Newell (2009), this might be because the classification question raised the salience of the target category, thus leading some participants to primarily base their judgments on exemplars from this category (the single-category conjunction users).

Our results extend Papadopoulos et al.'s (2011) finding of a majority preference for multiple-category feature conjunction to cases in which people had the option of making category-based predictions. Moreover, we found that those using different forms of feature conjunction reasoning showed characteristic patterns of visual attention during induction. For example, the significant minority who relied on single-category conjunction showed a stronger bias toward exemplars in the target category than did those who used multiple-category conjunction.

An important difference between the present design and Murphy and Ross (2010), in which single-category conjunction was found to dominate, was that the category membership of test instances in our inductive problems was less certain. For example, in Fig. 1, the probability that the test instance (triangle) belonged to the target category was .5. The corresponding probability in the Murphy and Ross

(2010) studies was .83 or more. An increase in the certainty of target category membership would mean that inclusion of information from other categories would generally have less impact on feature predictions. This might have contributed to the high rates of single-category conjunction reported by Murphy and Ross (2010).

We have suggested that multiple-category feature conjunction involves examining co-occurrences between given and to-be-predicted features across available exemplars, disregarding category bounds. There may be alternative explanations of our multiple-conjunction findings, however, that involve attending to category membership. For example, a reviewer helpfully observed that, in Fig. 1, when the target item was a triangle, participants might have noticed that three girls drew triangles, and that two of the three girls (Patricia and Mary) exclusively drew green triangles. They might then reason that if the target triangle was produced by Patricia or Mary it would be green (the multiple-conjunction response option). This alternative strategy considers only the exemplars with the given feature across relevant categories but inappropriately weights each category equally in the induction process (at least, it ignores likelihoods when weighting the contribution of each category). Because this strategy is sensitive to feature correlations across all relevant categories, it can be considered a variant of the multiple-conjunction strategy, although it does consider the exemplars in reference to their categories. Indeed, as discussed above, the eyetracking data suggest that people were sensitive to feature correlations and also to category membership, which is supportive of accounts such as this one, that integrate both forms of information. Future research using specifically tailored category structures will be able to discriminate these variants.

Perhaps the most important future goal will be to clarify the factors that drive people toward or away from feature-based reasoning strategies. Our use of a decision-only paradigm meant that people could examine the features of all exemplars while making inductive predictions. While there are analogues outside the laboratory (e.g., a physician may consult records of previous patients when making predictions about a novel case), clearly many inductive predictions are based on memories of exemplars and their category assignments. Although there is some evidence that feature conjunctions can influence inferences that are based solely on retrieval of previously learned categories (Newell et al., 2010), the extent of this influence has yet to be thoroughly examined.

Author Note This work was supported Australian Research Council Discovery Grant DP0770292 to the second and third authors.

We would like to thank Helen Paton for her assistance with programming and data collection.

Appendix

The mean proportion of dwell time per type of exemplar was calculated as shown in Eq 1. The total fixation duration for each exemplar is denoted $fixation_{i,j,k}$. Each exemplar is identified by three values: Type of exemplar is denoted by j , individual exemplars within each type are denoted by i (n exemplars per type), and the time bin is denoted k . The term $fixation_{misc,k}$ refers to the miscellaneous fixations within the stimulus presentation area in each time bin k (e.g., to category labels).

$$\text{Dwell time}_j = \frac{\sum_{i=1}^n fixation_{i,j,k}}{n \cdot \sum_{k=1}^3 \left(\left(\sum_{j=1}^5 \sum_{i=1}^n fixation_{i,j,k} \right) + fixation_{misc,k} \right)} \quad (1)$$

References

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409–429. doi:10.1037/0033-295X.98.3.409
- Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research*, 36, 1827–1837. doi:10.1016/0042-6989(95)00294-4
- Hayes, B. K., & Newell, B. R. (2009). Induction with uncertain categories: When do people consider the category alternatives? *Memory & Cognition*, 37, 730–743. doi:10.3758/MC.37.6.730
- Hayes, B. K., Heit, E., & Swendsen, H. (2010). Inductive reasoning. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1, 278–292. doi:10.1002/wcs.44
- Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, 27, 148–193. doi:10.1006/cogp.1994.1015
- Murphy, G. L., & Ross, B. H. (2007). Use of single or multiple categories in category-based induction. In A. Feeney & E. Heit (Eds.), *Inductive reasoning: Experimental, developmental, and computational approaches* (pp. 205–225). New York: Cambridge University Press.
- Murphy, G. L., & Ross, B. H. (2010). Category vs. object knowledge in category-based induction. *Journal of Memory and Language*, 63, 1–17. doi:10.1016/j.jml.2009.12.002
- Newell, B. R., Paton, H., Hayes, B. K., & Griffiths, O. (2010). Speeded induction under uncertainty: The influence of multiple categories and feature conjunctions. *Psychonomic Bulletin & Review*, 17, 869–874. doi:10.3758/PBR.17.6.869
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57. doi:10.1037/0096-3445.115.1.39
- Papadopoulos, C., Hayes, B. K., & Newell, B. R. (2011). Non-categorical approaches to feature prediction with uncertain categories. *Memory & Cognition*, 39, 304–318. doi:10.3758/s13421-010-0009-4
- 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. doi:10.1037/0278-7393.22.3.736