

# Speeded induction under uncertainty: The influence of multiple categories and feature conjunctions

BEN R. NEWELL, HELEN PATON, BRETT K. HAYES, AND OREN GRIFFITHS  
*University of New South Wales, Sydney, New South Wales, Australia*

When people are uncertain about the category membership of an item (e.g., Is it a dog or a dingo?), research shows that they tend to rely only on the dominant or most likely category when making inductions (e.g., How likely is it to befriend me?). An exception has been reported using speeded induction judgments where participants appeared to use information from multiple categories to make inductions (Verde, Murphy, & Ross, 2005). In two speeded induction studies, we found that participants tended to rely on the frequency with which features co-occurred when making feature predictions, independently of category membership. This pattern held whether categories were considered *implicitly* (Experiment 1) or *explicitly* (Experiment 2) prior to feature induction. The results converge with other recent work suggesting that people often rely on feature conjunction information, rather than category boundaries, when making inductions under uncertainty.

One of the central functions of categories is property induction (Anderson, 1991). Induction is relatively straightforward when an object's category is known with certainty. If we know that an object is a dog, we can be confident that it will bark, wag its tail, chase cats, and so on. However, induction becomes more complex when we are *uncertain* about category membership. If we glimpse an animal running through the bush, we may not be sure whether it is a dog or a dingo, and subsequent inductions (e.g., How likely is it to befriend me?) will depend on how we handle this uncertainty.

According to rational models of reasoning (e.g., Anderson, 1991), we should make inductive predictions using a *multiple-category approach*, whereby predictions are conditionalized on *all* possible categorizations of an object. Thus, if one is 70% certain that the animal is a dog and 30% certain that it is a dingo, one's induction about the chances of friendship should be *lower* than if one is 90% certain that it is a dog.

However, a large body of empirical work suggests that people do not incorporate information from multiple categories in induction under uncertainty. Instead, people tend to rely *only* on the most dominant or likely category. This leads people to rate the probability of unknown features to be equal even when they acknowledge different amounts of uncertainty in the initial categorization of an object. That is, both 70% and 90% certainties that an object was a dog lead to the same probability rating about the chance of friendship. This *single-category approach* to induction has been demonstrated in a variety of artificial and real-world categories (Hayes & Newell, 2009; Murphy & Ross, 1994, 2005; Ross & Murphy, 1996).

## Speeded Induction

A notable exception to the dominance of the single-category approach was reported by Verde, Murphy, and Ross (2005; hereafter, VMR). VMR speculated that a single-category approach is often attractive because it can be implemented without recourse to the complex Bayesian computations necessitated by the multiple-category approach (e.g., Anderson, 1991). Moreover, they pointed out that in many situations, the inclusion of information from other categories has a relatively small impact on the specific probability of an unknown feature; thus, a small improvement in the accuracy of a numerical prediction might be insufficient to promote the more complex multiple-category strategy (Verde et al., 2005).

However, VMR argued that changing the nature of the required prediction might boost multiple-category use. They made two changes to the standard paradigm. First, they asked for speeded judgments, rather than self-paced judgments. Second, they elicited binary feature predictions (e.g., Is the object likely to befriend you: yes or no?), rather than precise numerical probabilities. They argued that speeded judgments have been shown to rely on intuitive and automatic forms of reasoning and that binary judgments can be made more quickly and intuitively than the Bayesian analysis required to justify numerical predictions. According to VMR, both of these manipulations should lead participants to draw on information from multiple categories. Two experiments supported this interpretation: Speeded, binary feature inductions appeared to be based on information drawn from all possible categorizations of a given object, not just the dominant one.

---

B. R. Newell, ben.newell@unsw.edu.au

These results are important for two reasons. First, they stand in stark contrast to the dominance of single-category reasoning across a wide variety of experimental studies, and second, they suggest a differential influence of “consciously” and “unconsciously” considered categories. In the standard paradigm, participants are encouraged to reflect consciously on category membership prior to feature induction, and this appears to lead to a focus on the dominant category. In contrast, the speeded judgment, in which categories may be considered in some automatic or “unconscious” manner, appears to lead to integration of information from multiple categories (Murphy & Ross, 2007; Verde et al., 2005).

Given the important implications of these results, we sought in the present study to reexamine the VMR paradigm. Our primary goal was not simple replication (although we did replicate VMR’s basic effect) but the examination of an alternative explanation, precipitated by recent findings in related studies of category-based induction under uncertainty. Next, we describe the VMR study in detail and then present our alternative explanation, which was tested in the experiments that follow.

### The Verde et al. (2005) Study

VMR used a version of the children’s geometric drawings task developed by Murphy and Ross (1994). In the task, there were four categories labeled with children’s names: Bob, John, Ed, and Sam. Each child was said to have drawn 16 exemplars of geometric shapes, and some children were said to have preferences for different colors and shapes.

For example, Bob liked to draw triangles and color his shapes green; thus, he was the dominant category for the features *green* and *triangle*, because they were most likely (60% of the time) to have been drawn by him. John preferred *red* and *squares*, and objects with these features were most likely (60% of the time) to be drawn by him. For each critical feature pair (*green–triangle* and *red–square*), the distribution within the respective dominant (most likely) category (*Bob, John*) was equal: The probability that a green drawing would be a triangle (and vice versa) was 75%, and the probability that a red drawing would be square (and vice versa) was 75%. Thus, an induction about the missing feature (e.g., shape) of a green or red object should not differ if predictions are based solely on the dominant categories.

However, the distribution of critical features differed outside the dominant categories. Specifically, two children other than Bob drew green triangles, but no child other than John drew red squares. Thus, if information from outside the dominant category was used in induction, a green drawing should be rated as more likely to be a triangle than that a red drawing is likely to be square.

VMR trained participants to learn the category structures in a trial-by-trial paradigm (see the present Method section for details). In the induction test, participants were shown the phrase, “A picture that is:” followed by a feature word (e.g., *triangle*), and then, after a delay, were shown the phrase “is most likely to be:” followed by a second

feature word, which was a color if the given feature was a shape and a shape if the given feature was a color. The participants had to respond *yes* or *no* as quickly and accurately as possible.

The key finding was that participants responded faster to critical conjunctions that had occurred in more than one category (e.g., *green triangles*) than to the conjunctions that had appeared in only one (e.g., *red squares*). This faster responding could not, according to VMR, be attributed to single-category responding, because the distribution of conjunctions within the dominant category was identical. Thus, the speeded inductions were deemed to have been influenced by information from “unconsciously” considered multiple categories (Murphy & Ross, 2007; Verde et al., 2005).

### A Role for Feature Conjunctions

Although this explanation is consistent with VMR’s data, an alternative account can be advanced. Crucially, the number of feature co-occurrences and the number of categories in which objects appeared were confounded in VMR’s design. Green triangles did appear in more *categories* than did red squares but they also appeared *more often* across all training items. Thus, as VMR acknowledged, it is impossible to determine whether the decreased response times (RTs) were due simply to greater frequency of occurrence or to the integration of information across categories. The former interpretation is perhaps more likely, given that a simple extrapolation of the VMR result leads to the counterintuitive prediction that responding should be further facilitated by increasing the number of categories in which an object occurs.

Determining what is responsible for the VMR result is crucial, because it speaks directly to whether participants use any kind of *categorical representation* when making inductions or simply rely on the presence of conjunctions in individual exemplars. This issue is especially pertinent given several recent findings implicating the use of feature conjunction reasoning in the standard induction-under-uncertainty paradigm (Hayes, Ruthven, & Newell, 2007; Murphy & Ross, 2010; Papadopoulos, Hayes, & Newell, in press). These studies showed that when people are uncertain about an object’s category membership, they default to inductive predictions based on feature correlations, often eschewing category boundaries. So, in our earlier example, people could make a prediction about the friendliness of the unknown animal by simply retrieving associations between known features (e.g., animal of a certain size and shape) and likelihood of friendliness, without a preliminary stage of categorization.

The aim of the present experiments was to distinguish between the role played by category membership and the role played by feature conjunctions, by manipulating these factors orthogonally. In order to examine the possible influence of implicit and explicit consideration of categories, in Experiment 1, categories were not mentioned prior to feature induction (as per VMR), but in Experiment 2, an explicit category membership judgment was required immediately before the feature induction was made.

## EXPERIMENT 1 Multiple Categories Versus Feature Conjunctions

The simplest way to examine the effects of category membership and feature conjunction is to use a factorial design with two levels of each factor—for example, membership (few vs. many categories) and feature conjunction presentations (few vs. many). The feature conjunction account predicts a main effect of number of presentations on RT (*many* faster than *few*) but no effect of category membership; the simple extrapolation of VMR's result leads to the prediction that responding will be faster to items in more categories.

Following VMR, we operationalized *many* feature conjunction presentations as 13 and *few* as 9. For category membership, we compared items that were members of one or two categories with those that were members of four categories. We chose one or two categories (rather than a simple comparison of two vs. four categories) in order to preserve the VMR stimulus construction in which the critical evidence came from finding shorter RTs for items presented 13 times in two categories than for items presented 9 times in one category.

### Method

**Participants.** Thirty-two first-year psychology students from the University of New South Wales participated for course credit.

**Materials and Design.** The structure of the training categories is displayed in Table 1. Four categories (*Bob*, *John*, *Ed*, and *Sam*) were created with 20 exemplars, each consisting of two features: shape and color. Critical conjunction pairs appeared in four categories for a total of 13 presentations (e.g., *red-square*: three times each in the categories *Bob*, *Ed*, and *Sam* and four times in the category *John*), in four categories for a total of 9 presentations (e.g., *yellow-diamond*: twice each in *Bob*, *John*, and *Ed* and three

times in *Sam*), in two categories for a total of 13 presentations (e.g., *green-triangle*: 9 times in *Bob* and 4 in *Sam*), or in one category for a total of 9 presentations (e.g., *blue-circle*: 9 times in *Ed*).

The key VMR comparison was preserved by ensuring that distributions of the critical pairs were the same within the two dominant categories *Bob* and *Ed*: In both categories, a drawing was 60% likely to have each critical feature (e.g., *green* or *triangle* in *Bob*'s case, *blue* or *circle* in *Ed*'s), and a drawing with one critical feature was 75% likely to have the other. Crucially, the conditions differed in the distribution of critical feature pairs in the nondominant categories, because *green triangles* were also drawn by *Sam*, but only *Ed* drew *blue circles*.

**Procedure.** The participants were given a study booklet containing the 20 drawings of each child on four separate pages. They were told to familiarize themselves with each child's drawings. After 10 min, the booklets were removed, and the learning phase was conducted on a computer. In the learning phase, all 80 exemplars from the four categories were presented in random order. On each test trial, the question "Who drew this picture?" was displayed at the top of the screen. After 1 sec, 1 exemplar and the four response options (*Bob*, *John*, *Sam*, *Ed*) were displayed. The participants were asked to respond as quickly and accurately as possible by pressing an assigned key. If a picture could have been drawn by more than one person, they were told to give the best answer. Feedback was displayed and remained on screen for 2 sec for correct and 4 sec for incorrect responses (to encourage accurate responding). After the first learning phase, the participants were allowed to consult the study booklet for a further 5 min before completing a second, identical learning phase.

The final speeded induction test consisted of 8 practice trials followed by 72 test trials. Each trial started with a blank screen (for 1,500 msec); then "A picture that is" appeared on the left-hand side of the screen. After a 500-msec delay, a feature word (e.g., *green*) appeared in the middle of the screen, along with "is most likely to be." After a further 500-msec delay, a second feature word (e.g., *triangle*) appeared below the first, and a *yes* or *no* response was elicited. Each of the four critical feature pairs was presented 8 times, for a total of 32 trials. The remaining 40 trials comprised 16 in which a critical feature was paired with a filler feature (e.g., *green-oval*) and 24 in which each filler was presented with all

Table 1  
Category Structure Used in Experiments 1 and 2

Exemplar	Category 1 (Bob)		Category 2 (John)		Category 3 (Ed)		Category 4 (Sam)	
	Color	Shape	Color	Shape	Color	Shape	Color	Shape
1	<b>green</b>	<b>triangle</b>	<b>red</b>	<b>square</b>	<b>blue</b>	<b>circle</b>	<b>yellow</b>	<b>diamond</b>
2	<b>green</b>	<b>triangle</b>	<b>red</b>	<b>square</b>	<b>blue</b>	<b>circle</b>	<b>yellow</b>	<b>diamond</b>
3	<b>green</b>	<b>triangle</b>	<b>red</b>	<b>square</b>	<b>blue</b>	<b>circle</b>	<b>yellow</b>	<b>diamond</b>
4	<b>green</b>	<b>triangle</b>	<b>red</b>	<b>square</b>	<b>blue</b>	<b>circle</b>	yellow	oval
5	<b>green</b>	<b>triangle</b>	red	rectangle	<b>blue</b>	<b>circle</b>	yellow	oval
6	<b>green</b>	<b>triangle</b>	red	rectangle	<b>blue</b>	<b>circle</b>	yellow	oval
7	<b>green</b>	<b>triangle</b>	red	rectangle	<b>blue</b>	<b>circle</b>	orange	diamond
8	<b>green</b>	<b>triangle</b>	white	square	<b>blue</b>	<b>circle</b>	orange	diamond
9	<b>green</b>	<b>triangle</b>	white	square	<b>blue</b>	<b>circle</b>	orange	diamond
10	green	rectangle	white	square	blue	oval	<b>red</b>	<b>square</b>
11	green	rectangle	yellow	circle	blue	oval	<b>red</b>	<b>square</b>
12	green	rectangle	yellow	circle	blue	oval	<b>red</b>	<b>square</b>
13	orange	triangle	yellow	circle	white	circle	white	oval
14	orange	triangle	yellow	circle	white	circle	white	oval
15	orange	triangle	blue	diamond	white	circle	orange	rectangle
16	<b>yellow</b>	<b>diamond</b>	blue	diamond	<b>yellow</b>	<b>diamond</b>	orange	rectangle
17	<b>yellow</b>	<b>diamond</b>	blue	diamond	<b>yellow</b>	<b>diamond</b>	<b>green</b>	<b>triangle</b>
18	<b>red</b>	<b>square</b>	blue	diamond	<b>red</b>	<b>square</b>	<b>green</b>	<b>triangle</b>
19	<b>red</b>	<b>square</b>	<b>yellow</b>	<b>diamond</b>	<b>red</b>	<b>square</b>	<b>green</b>	<b>triangle</b>
20	<b>red</b>	<b>square</b>	<b>yellow</b>	<b>diamond</b>	<b>red</b>	<b>square</b>	<b>green</b>	<b>triangle</b>

Note—Critical conjunction items are in bold. In Experiment 2, the assignment of features (shapes, colors) was counterbalanced across categories.

other possible features (e.g., *orange–diamond*). *Yes* and *no* were the correct responses on an equal number of trials.

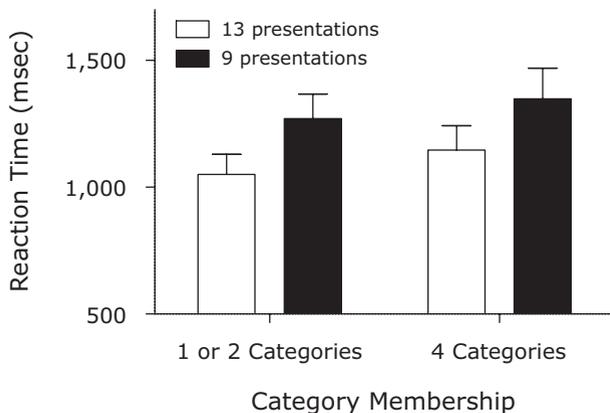
## Results

**Learning phase data.** Accuracy in the second learning phase was 61.45%, significantly above chance (25%) [ $t(31) = 32.10, p < .001$ ], indicating that the participants were able to learn this relatively complex categorization task.

**Speeded induction test data.** The participants were very accurate, with 88% to 94% correct inductions about the features of the four critical pairs. Accuracy did not vary as a function of category membership or number of presentations (largest  $F < 1.5$ ).<sup>1</sup>

The RT data for correct feature inductions are shown in Figure 1.<sup>2</sup> For the primary analyses, we collapsed across the one versus two category membership items, so that we could examine the main effect of few (one or two) versus many (four) categories. There was a main effect of number of presentations, with items presented 13 times (1,098 msec) responded to faster than those presented 9 times (1,309 msec) [ $F(1,31) = 5.50, p = .026$ ]. In contrast, there was no significant main effect of category membership [ $F(1,31) = 1.63, p = .21$ ], although the trend was opposite to that predicted by the extrapolation of the VMR result, with faster responses for items presented in one or two categories (1,160 msec) than for those presented in four categories (1,247 msec). The interaction between number of presentations and category membership was not significant ( $F < 1$ ).

Simple effects analyses were conducted to explore the specific VMR finding.<sup>3</sup> These revealed that items presented 13 times in two categories were responded to faster than those presented 9 times in one [ $F(1,31) = 7.37, p = .031$ ], despite their having identical distributions within their dominant categories. The difference in RTs to items presented 13 times in four categories and 9 times in four did not reach significance [ $F(1,31) = 1.88, p = .18$ ], but Figure 1 shows a trend for faster responding to the more frequently exposed items.



**Figure 1.** Mean reaction times to the critical conjunction items in the speeded induction test in Experiment 1. Conjunction items were presented during training either 9 or 13 times and were members either of one or two or of four categories. See the text for further clarification of the design. Error bars represent *SEM*.

## Discussion

Experiment 1 replicated VMR's main finding, showing that feature conjunctions that were equated in their dominant categories (e.g., *green triangles* and *blue circles*) were responded to differently in speeded induction judgments. However, the faster responding to *green triangles* appeared to be due to frequency of exposure of the conjunction of features, rather than to the number of categories these exemplars were shown in. This is because at the overall level, the frequency of exemplar presentation significantly affected RT, but the number of categories in which an exemplar was shown did not. In fact, contrary to the simple extrapolation of the VMR result, increasing the number of categories with a given feature tended to slow induction (see Figure 1).

## EXPERIMENT 2

### Explicit Consideration of Categories

VMR speculated that the multiple-category responding they observed was due to automatic or implicit consideration of information from multiple categories (see also Murphy & Ross, 2007). Experiment 1 challenged the multiple-category account of VMR's results by demonstrating that frequency of feature co-occurrence, and not category membership, drives the facilitation in inductive responding. Experiment 2 examined whether this pattern would hold when categories were explicitly considered (or where consideration was at least encouraged) immediately before feature induction. This was done by preceding the feature induction question with a binary category membership judgment.

According to VMR's characterization, the explicit consideration of categories should push responding toward single-category reasoning. This would result in no RT differences for critical feature conjunctions that are equated in their dominant categories. In contrast, the feature conjunction account, in which category membership is ignored, predicts no effect of making categories more salient and, thus, a data pattern similar to that in Experiment 1.

## Method

**Participants.** Sixty-four first-year psychology students from the University of New South Wales participated for course credit.

**Materials and Design.** The category structure was the same as that in Experiment 1, but we added a counterbalanced version of the stimuli. Thirty-one participants completed a version with the assignment of features shown in Table 1, and 33 completed a counterbalanced version in which colors and shapes were randomly paired and then exchanged across the category structure (e.g., all red shapes in the original version became green shapes in the counterbalanced version, and vice versa).

**Procedure.** The procedure was identical to that in Experiment 1, with the addition of a category membership question in the speeded induction test. On each test trial, the participants were shown a "picture that is," and then, after a 500-msec delay, a feature word was shown and they were asked "was most likely drawn by," followed by two children's names. For each trial, the two names consisted of one from the most likely category and one from an unlikely category. Immediately following a response to the category question, the feature prediction question was asked. There was a 1-sec intertrial interval.

## Results

**Learning phase data.** Accuracy in the second learning phase was 56.46%, significantly above chance (25%) [ $t(63) = 26.47, p < .001$ ].

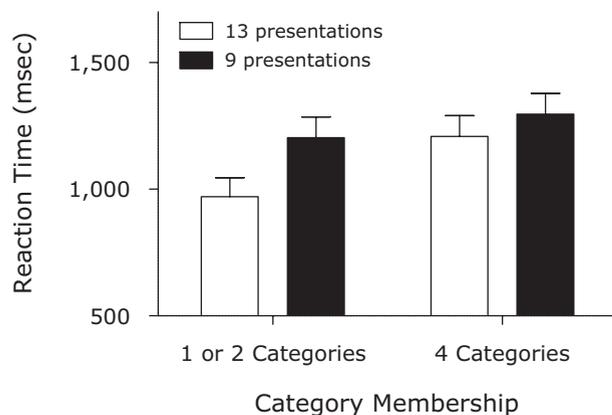
**Speeded induction test data.** Accuracy of feature inductions was high overall and was affected by category membership, but not by number of presentations. Predictions about items that appeared in one or two categories (82%) were more accurate than those for items presented in four categories (64%) [ $F(1,62) = 64.28, p < .001$ ].

For the primary RT analyses (Figure 2), we again collapsed across the one versus two category membership items so that we could examine the main effect of few (one or two) versus many (four) categories. In addition, we examined RTs only for judgments on which both feature induction and category membership judgments were correct.<sup>4</sup> Consistent with Experiment 1, conjunctions presented 13 times were responded to faster (1,089 msec) than those presented only 9 times (1,250 msec) [ $F(1,63) = 10.62, p = .002$ ]. In addition, responses to items appearing in one or two categories (1,087 msec) were significantly faster than responses to those appearing in four categories (1,252 msec) [ $F(1,63) = 9.52, p = .003$ ]. The interaction between the factors was not significant [ $F(1,63) = 2.46, p = .12$ ], although a trend is apparent suggesting a larger difference between exposure frequencies in the one- or two-category conditions.

Simple-effects analyses (see note 3) revealed that items presented 13 times in two categories were responded to faster than those presented 9 times in one category [ $F(1,63) = 12.09, p < .001$ ], again replicating VMR. The other simple effect was not significant, but the trend was for faster responding to the more frequently presented item [ $F(1,63) = 1.68, p = .200$ ].

## Discussion

Consistent with Experiment 1, critical conjunctions presented more frequently during training were responded to faster at test, independently of category membership.



**Figure 2.** Mean reaction times to the critical conjunction items in the speeded induction test in Experiment 2. Conjunction items were presented during training either 9 or 13 times and were members either of one or two or of four categories. See the text for further clarification of the design. Error bars represent SEM.

These data seem most readily explained by the noncategorical feature conjunction account. Furthermore, rather than leading participants toward single-category reasoning, explicit categorization of test items appeared to interfere with subsequent induction. Items that appeared in four categories were responded to less accurately and more slowly than were those appearing in one or two categories.

## GENERAL DISCUSSION

The present experiments show that membership in multiple uncertain categories does not facilitate inductive judgments but that additional pairings of given and predicted features do. That is, the frequent presentation of exemplars that consist of a conjunction of two features facilitates subsequent prediction of one feature from the other.

These data challenge the claim that VMR's basic effect was due to the use of multiple-category reasoning in feature induction. Our data are consistent with a feature conjunction approach to uncertain induction, whereby people make feature predictions on the basis of what they know about other exemplar features, rather than on the basis of category membership. This is an important result, because most previous work on induction has assumed that categories are central to induction. Our data suggest, however, that when category membership is uncertain, people often favor predictions based on feature conjunction (cf. Murphy & Ross, 2010). The processing mechanisms involved in feature conjunction have yet to be fully specified, but the present data are consistent with exemplar-based accounts (e.g., Nosofsky & Palmeri, 1997), in which the given feature (e.g., triangle) acts as a retrieval cue for exemplars in memory with this feature. Critical features that have occurred together more often in training (e.g., green and triangle) are then more readily and accurately recalled at test, due to stronger associations in memory between those features.

The data do not support the claim that categories considered *implicitly* (Experiment 1) or *explicitly* (Experiment 2) result in different patterns of performance in induction under uncertainty. The principal difference between the experiments—the shift toward even slower responding to items spaced across multiple categories in Experiment 2—suggests that when feature conjunction information is used as the default, membership in multiple categories has a deleterious effect on responding. One caveat to this conclusion is that in Experiment 2, participants responded slightly more accurately (84.5% vs. 80.1%) and significantly faster to items seen 13 times in *two* categories than to those seen 9 times in *one*. However, the overall pattern in the data of both experiments suggests that the improved performance in induction was due to the increase in exposures to conjunctions, and not to multiple-category membership.

Taken together, these data suggest a convergence between speeded induction tasks (the present experiments) and more “deliberative” feature prediction studies (Hayes et al., 2007; Murphy & Ross, 2010; Papadopoulos et al.,

in press). VMR contrasted their results in the speeded task with behavior in another experiment that used a self-paced, numerical prediction task but the same category structure. They interpreted the results of the self-paced task as indicating single-category reasoning, thereby suggesting that different processes are invoked as a function of task requirements. However, the pattern of responding in their self-paced task *could* also be the result of feature conjunction reasoning, because within each dominant category, the feature predicted by single-category reasoning was always the one that was paired most frequently with the feature given during the induction test. This suggests a common mechanism for speeded and self-paced induction. Indeed, many of the results that have previously been assumed to show reliance on category-based reasoning (single or multiple) may reflect feature conjunction reasoning (e.g., Murphy & Ross, 1994, Experiments 1 and 2), suggesting that people often prefer to make inductive predictions on the basis of the certain or known properties of an object, rather than on the basis of uncertain category membership.

#### AUTHOR NOTE

The support of the Australian Research Council (Grant DP: 0770292, awarded to the first and third authors) is gratefully acknowledged. Correspondence should be addressed to B. R. Newell, School of Psychology, University of New South Wales, Sydney, NSW 2052, Australia (e-mail: ben.newell@unsw.edu.au).

*Note*—Accepted by Cathleen M. Moore's editorial team.

#### REFERENCES

- ANDERSON, J. (1991). The adaptive nature of human categorization. *Psychological Review*, **98**, 409-429. doi:10.1037/0033-295X.98.3.409
- HAYES, B. K., & NEWELL, B. R. (2009). Induction with uncertain categories: When do people consider the alternative categories? *Memory & Cognition*, **37**, 730-743. doi:10.3758/MC.37.6.730
- HAYES, B. K., RUTHVEN, C., & NEWELL, B. R. (2007). Inferring properties when categorization is uncertain: A feature conjunction account. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Conference of the Cognitive Science Society* (pp. 209-214). Mahwah, NJ: Erlbaum.
- MURPHY, G. L., & ROSS, B. H. (1994). Predictions from uncertain categories. *Cognitive Psychology*, **27**, 148-193. doi:10.1006/cogp.1994.1015
- MURPHY, G. L., & ROSS, B. H. (2005). The two faces of typicality in category-based induction. *Cognition*, **95**, 175-200. doi:10.1016/j.cognition.2004.01.009
- 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). Cambridge: Cambridge University Press.
- MURPHY, G. L., & ROSS, B. H. (2010). Category vs. object knowledge in category-based induction. *Journal of Memory & Language*, **63**, 1-17. doi:10.1016/j.jml.2009.12.002
- NOSOFSKY, R. M., & PALMERI, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, **104**, 266-300. doi:10.1037/0033-295X.104.2.266
- PAPADOPOULUS, C., HAYES, B. K., & NEWELL, B. R. (in press). Non-categorical approaches to induction with uncertain categories. *Memory & Cognition*. doi:10.3578/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, & Cognition*, **22**, 736-753. doi:10.1037/0278-7393.22.3.736
- 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.

#### NOTES

1. Performance on the noncritical pairs was not of direct theoretical interest, but analyses revealed high levels of accuracy (e.g., average accuracy on *critical-filler* pairs was 88% in Experiment 1 and 82.9% in Experiment 2).
2. RTs were trimmed; values two *SDs* above or below the mean were replaced by the cutoff value for each participant (5.2% of the induction data in Experiment 1 and 4.9% in Experiment 2).
3. The interaction was not significant, thus technically challenging the legitimacy of these simple-effect analyses. However, we felt that the importance of examining the basic VMR effect (i.e., establishing whether it could be replicated) justified conducting them.
4. We were interested in trials on which participants made correct category judgments *and* correct feature inductions. An analysis of RTs including trials on which categorization errors were made but feature inductions were correct did not change the conclusions of Experiment 2.

(Manuscript received February 2, 2010;  
revision accepted for publication June 1, 2010.)