Journal of Experimental Psychology:
Learning, Memory, and Cognition
2016, Vol. 42, No. 7, 1050–1067

© 2016 American Psychological Association
0278-7393/16/$12.00 http://dx.doi.org/10.1037/xlm0000222

Eyetracking Reveals Multiple-Category Use in Induction

Stephanie Y. Chen
New York University

Brian H. Ross
University of Illinois at Urbana-Champaign

Gregory L. Murphy
New York University

Category information is used to predict properties of new category members. When categorization is uncertain, people often rely on only one, most likely category to make predictions. Yet studies of perception and action often conclude that people combine multiple sources of information near-optimally. We present a perception-action analog of category-based induction using eye movements as a measure of prediction. The categories were objects of different shapes that moved in various directions. Experiment 1 found that people integrated information across categories in predicting object motion. The results of Experiment 2 suggest that the integration of information found in Experiment 1 were not a result of explicit strategies. Experiment 3 tested the role of explicit categorization, finding that making a categorization judgment, even an uncertain one, stopped people from using multiple categories in our eye-movement task. Experiment 4 found that induction was indeed based on category-level predictions rather than associations between object properties and directions.

Keywords: category-based induction, reasoning, implicit processes

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Navigating our everyday world involves making predictions about the properties of new objects we encounter. Is that food too spicy to eat? Is that dog on the street friendly or dangerous? Is that object going to shatter if it falls off the table? Such inductions often rely on category-level knowledge. If the object is a plate, you would likely predict that it will shatter easily. If the item is instead a Frisbee, you might predict that it will not shatter easily.

This problem becomes more complex when we are uncertain of what category an item belongs to. At first glance, we may not know for sure whether the thing on the table belongs to the category of plate or Frisbee. Normatively, under such uncertain circumstances we should base our predictions on information from all the object's possible categories weighted by how likely it is that the item belongs to that category. To decide if the item will shatter when dropped, multiply the probability that the object in question is a plate by the probability that a plate would shatter when dropped. Next multiply the probability that the object is a Frisbee by the probability that a Frisbee would shatter. The sum of the two products is the probability that the item will shatter when dropped (assuming these two categories are exhaustive and mutually exclusive). This is consistent with normative principles and Bayesian approaches in which people weight different possibilities by their prior likelihoods (see Anderson, 1991, for such an approach).

Research on category-based induction under uncertainty has asked if people are normative when making these types of predictions. In particular, it has been concerned with whether people use information from multiple categories (like plates and Frisbees) during induction. These studies have found that subjects' predictions surprisingly often do not reflect the normative integration of information across categories described above (we review this literature in more detail shortly). Instead, these predictions are often based on only the most likely category, disregarding relevant information from less likely alternatives (Hayes & Chen, 2008; Hayes & Newell, 2009; Malt, Ross, & Murphy, 1995; Murphy, Chen, & Ross, 2012; Murphy & Ross, 1994; Ross & Murphy, 1996; Verde, Murphy, & Ross, & 2005). However, under other circumstances, people may use multiple categories, depending on the question asked or details of the stimulus presentation and category structure (Griffiths, Hayes, Newell, & Papadopoulos, 2011; Murphy & Ross, 2010b; Papadopoulos, Hayes, & Newell, 2011; Verde et al., 2005). This investigation attempts to further clarify when people do and do not use multiple categories by contrasting traditional category-based induction tasks with those from other areas of cognitive science in which people seem to be able to integrate information in a normative manner.

Research from perception and motor control has consistently found that people are able to normatively integrate information from multiple sources, outcomes, or possibilities (Ernst & Banks, 2002; Haruno, Wolpert, & Kawato, 2001; Kersten, Mamassian, & Yuille, 2004; Tassinari, Hudson, & Landy, 2006; Trommershäuser, Landy, & Maloney, 2006; Trommershäuser, Maloney, &

This article was published Online First January 4, 2016.

Stephanie Y. Chen, Department of Psychology, New York University; Brian H. Ross, Department of Psychology, University of Illinois at Urbana-Champaign; Gregory L. Murphy, Department of Psychology, New York University.

We thank Michael Grubb, and the NYU ConCats Research Group for helpful comments, and Rebecca Bainbridge for indispensible help in data collection and technical assistance. The research was supported in part by National Science Foundation Grants BCS-1128769 and 1128029.

Correspondence concerning this article should be addressed to Stephanie Y. Chen, who is now at the University of Chicago Booth School of Business, 5807 S. Woodlawn Avenue, Chicago, IL 60637. E-mail: stephanie.chen@chicagobooth.edu

Landy, 2008). In perception, Bayesian models are used to explain how the visual system takes ambiguous inputs and returns the most likely percept. For example, people are able to integrate different types of sensory information (e.g., haptic, visual) in a manner that optimally accounts for their uncertainty (variance in these observations; Ernst & Banks, 2002). In motor control, one action may be best suited to achieve a goal, given the state of the world. However, because perception is not perfect, the state of the world is uncertain. Models of action propose that people integrate information about the likelihood of the possible states of the world to make near-optimal actions (Haruno et al., 2001). Trommershäuser, Körding, and Landy (2011) summarize many examples of such integration of options in the perceptual domain.

It is surprising that such work on perception and action, which often involves making a prediction from uncertain cues, so often reveals near-normative use of those cues, when people often do not combine two categories in simple tasks of category-based induction. One difference between these two domains is that the category-based induction experiments often allow or encourage the use of focusing strategies, in which people choose to attend to only one category. Some people may believe that this strategy is actually normative (Murphy et al., 2012). The perception-action experiments often do not ask for predictions but rather require motor responses. In a number of cases, people may not be aware of the underlying perceptual cues; in other cases, responses are speeded and do not permit slower strategies that are apparently used in the induction tasks (e.g., categorizing the object as most likely a plate and then basing your prediction on only the probability that a plate will shatter).

We first review the evidence for different types of reasoning in the traditional induction task. We then consider induction tasks that are more like the perception-action tasks cited above and have consistently yielded multiple-category use. Our investigation aims to further understand why these tasks avoid the single-category focus found in category-based induction to answer additional questions about what determines whether induction uses multiple or single categories. To do this we have created a new induction task (similar to tasks used in perception-action research) based on our earlier study in which people had to catch moving shapes (Chen et al., 2014).

How do People Use Categories in Induction?

Initial research in category-based induction consistently found that most people tend to base their predictions on only a single category when making predictions under uncertainty. This finding has been shown with both real and artificial categories. The single-category focus occurs even when people have acknowledged that they are not sure what category an item is in just before making their induction (Hayes & Newell, 2009; Murphy & Ross, 1994; Ross & Murphy, 1996; Verde et al., 2005), when the possible categories are equally likely and there is no rationale for selecting one category over another (Murphy et al., 2012, Experiment 3), and when people have used information from another category to answer a question just before making the induction (Ross & Murphy, 1996).

However, more recent research has found a more diverse set of induction strategies, depending on individual differences (Hayes & Chen, 2008; Murphy & Ross, 2010b) and procedural or stimulus variables. Although we have focused on differentiating induction strategies that use multiple versus single categories, another (potentially orthogonal) distinction is whether subjects use category-level or feature-based strategies to make their predictions. To illustrate we provide an example of each strategy that uses information from only a single category below.

In a standard paradigm (Murphy & Ross, 1994), subjects look at a display that shows examples of two to four categories. They typically are told a property of a new entity, asked a categorization question, and then make a prediction about another property. Figure 1 shows a simplified display, for illustrative purposes. The cover story for this task is often that children (categories) have used a computer program to draw shapes in different colors (that we represent as shadings). If subjects are told that a new item is a triangle and asked to predict its shading, they may focus on the category that has the most triangles, Nina, and predict the shading that is most frequent in the category, solid black; this would be a *category-level* strategy that uses information from only a single category. If instead subjects use a feature-based approach, they might base their predictions on only the items in the most likely category that have the mentioned feature (triangle). This strategy would lead them to predict the shading that is most frequent among triangles in the most likely category, dotted. This type of feature-based strategy is often referred to as a *feature conjunction strategy*. Many studies have found evidence that people use such a strategy in the types of category-based induction tasks described above (Griffiths et al., 2011; Murphy & Ross, 2010a; Newell et al., 2010; Papadopoulos et al., 2011).

The distinction between category-level and feature-conjunction strategies can also be made for induction strategies that use multiple categories. That is, people can either base their predictions on the probability of the different shadings for each category (weighted by how likely each category is), or on the frequency of shadings for only the items with the queried feature, essentially ignoring category membership—for example, predict the most common shading of all triangles drawn by both Nina and Lindsey. (In Figure 1, both strategies make the same prediction for the multiple categories case, but other studies, discussed below, have used more complex designs that can distinguish between these strategies.)

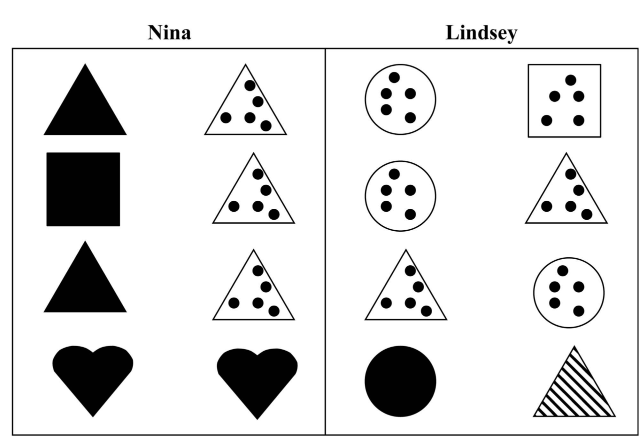

Figure 1. Example of stimuli used previous category-based induction studies.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Although it has been proposed that people tend to base their inductions on items from multiple categories when they can use a feature-conjunction strategy (Papadopoulos et al., 2011), single-category use makes up a significant portion of feature-based predictions. Thus, the question of how to promote multiple-category use in induction remains relevant. In a series of studies, Griffiths et al. (2011, 2012) found evidence for feature-based strategies based on both multiple and single categories. For example, Griffiths et al. (2011, Experiment 1) found that the vast majority of responses reflected a feature-based strategy. These feature-based responses, however, were not consistent in how many categories they were based on: Some were based on information from a single category and some on information from multiple categories. Furthermore, a number of studies have documented a preference to focus on single categories even when the feature-based strategy is taken into account. For example, Murphy et al. (2012, Experiment 1) found that, using the standard questions, people made predictions based on a single category two thirds of the time. An interesting finding was that when people were explicitly asked to select the best strategy to answer induction questions (Experiment 5), 18 subjects claimed that focusing on the most likely category was correct, 14 said that multiple categories should be used, but only 4 said that the categories should be ignored, and the given feature used. Thus, many people may believe that focusing on only the most likely category is the correct strategy.

Finally, it should be noted that studies done with natural categories generally do not permit a feature-based process. Unlike the studies described above where uncertainty regarding categorization was created by mentioning a feature that was represented in multiple categories, in studies done with natural categories, often only the categories are mentioned (e.g., "She thought the man was the realtor, but she wasn't sure and wondered if he might be the cable repairman."). These studies also have found evidence for single-category reasoning (Malt et al., 1995; Ross & Murphy, 1996; Zhu & Murphy, 2013).

Our own conclusion from this set of results is that people may use a number of different strategies to make predictions when a category is uncertain. Furthermore, the details of the question and display are important. For example, Griffiths et al. (2012, Experiment 2) found large differences in whether people used single or multiple categories depending on whether the categories were learned before test and (to a lesser degree) how coherent the categories were. Murphy et al. (2012, Experiment 1) found a doubling of multiple-category use when subjects were asked about the likelihood that the target item belonged to each of the possible categories before the induction question versus when they were only asked about the most likely category. Therefore, the scientific question is not whether people use single or multiple categories—they use both, depending on a number of variables. However, it is still of interest to understand why and in what situations single category strategies are used when someone is not certain that an object is in that category—particularly since in some situations, that is the dominant response. The present study furthers the investigation of why people sometimes focus inappropriately on a single category when categorization is uncertain. We designed our categories so that we could detect the use of single or multiple categories whether people used a feature-based or category-level strategy.

Implicit Responding and Multiple-Category Use

Our proposal is that people focus on a single category as part of a cognitive strategy Evans (2007) calls the Singularity Principle, that people generally only consider one hypothetical possibility at a time (related to Stanovich's, 2009, claim that people are cognitive misers). The perception-action studies described above often show integration of information across different cues because people do not explicitly identify the predictions of the different sources. For example, when a speeded motor response might fall into a reward or loss area (Trommershäuser et al., 2006; Trommershäuser, Maloney, & Landy, 2003), people do not explicitly identify the probabilities of these two outcomes. They have no time to do so. Instead, their behavior can be explained as an associative response that pushes them toward the positive area and away from the loss area, depending on their motor variability and the sizes of the rewards. We have drawn an analogy to this type of experiment to construct similar speeded tasks in which category-based induction might show use of multiple categories.

In a study closely related to the present experiments, Chen et al. (2014) examined the role of response mode on category use by contrasting a verbal induction task with a game-like motor induction task. Subjects had to predict the direction of moving shapes either verbally, or by catching fast-moving shapes with their cursor in a game-like task. Inductions in the motor task showed evidence of integration across categories, and, similar to previous work, verbal inductions showed no evidence of integration across categories. This pattern of results mirrors research in decision making that suggests that people are more optimal and are better able to normatively integrate multiple uncertain outcomes when executing speeded visuomotor tasks than when making the equivalent decision in a standard economic choice task (Trommershäuser et al., 2006, 2008). Interestingly, when people are required to first make categorizations in perceptual tasks involving noise (i.e., uncertainty), they then seem to ignore evidence that is not consistent with that categorization (Stocker & Simoncelli, 2007). However, a perception task is not induction, and this result did not use learned categories. Therefore, the tasks used in Chen et al. (2014) and in the present research complete the analogy by teaching people categories and then testing how they make inductions to new items.

In Chen et al. (2014), as here, we referred to the verbal predictions as *explicit* and the motor predictions (here, eyetracking) as *implicit*. However, we should clarify that our goal is not to draw conclusions about two different systems of reasoning (cf. Sloman, 1996) nor do we claim that our implicit tasks are completely outside of conscious awareness. Rather, we start from the observation that different kinds of behavior seem to be differentially sensitive to multiple categories and seek to explain why. We do not assume that all explicit inductions use the same strategy, for example, as the use of multiple categories depends on individual differences and procedural variables. Nor do we assume that all implicit inductions must be the same. Processes that are often cited as implicit have a number of properties that may lead to the consideration of multiple categories, but the details of each process must be known to make such a prediction. For the purposes of this article, we mean that the implicit inductions are not the subject's focus in doing the task and, thus, likely avoid some of the strategies used when subjects explicitly make such induction.

The Current Research

The present article has three main goals that all serve to further understand why the single-category focus occurs. First, as our claim is that the single-category focus arises from cognitive strategies that occur when people have the goal of making an induction, we build upon the Chen et al. (2014) experiments with a new eyetracking measure that is even less likely to involve explicit predictions of an object's properties (its direction of motion). The new task does not ask, directly or indirectly, about direction. Second, we explore an important variable in whether people will focus on a single category, namely initial categorization judgments. The differences between the implicit and explicit tasks from our studies suggest that implicit induction may allow for more normative use of category information because it avoids an initial categorization of test items. In explicit prediction tasks, subjects were often asked to categorize the shape before the induction (Chen et al., 2014; Hayes & Newell, 2009; Malt et al., 1995; among many others). In everyday life, people may spontaneously choose a most-likely category of an object, even if they are not sure about it. There is evidence that subjects overcommit to their initial categorization and make their predictions consistent with it—indeed, even if they arbitrarily choose a category, many people make a prediction based on that categorization (Lagnado & Shanks, 2003; Murphy et al., 2012). However, in implicit induction tasks, subjects are generally not encouraged to think about an item's categorization before induction (but see Newell et al., 2010; Stocker & Simoncelli, 2007), and the time pressure that these predictions are made under makes it unlikely that subjects categorized spontaneously. In Experiment 3, we investigate whether initial categorization influences induction even in our eyetracking task.

Third, we investigate whether people's multiple-category use when making implicit predictions is a result of category-level or feature-based information. That is, are they sensitive to the fact that objects in a category generally go in a certain direction (e.g., items in Category 2 go to the upper right), or are they only using associations between the object's other properties and direction (e.g., squares tend to go to the upper right)? This speaks to issues that have received much attention in the category-based induction literature as discussed above (e.g., Griffiths et al., 2011, 2012; Murphy & Ross, 2010a; Papadopoulos et al., 2011).

Overall, the results from four experiments provide evidence that induction involves integration of information across categories when explicit reasoning strategies are avoided. We demonstrate this by providing evidence that implicit predictions measured via eye movements show more normative integration of information from multiple categories (Experiment 1) and that subjects do not explicitly report using such a strategy (Experiment 2). In Experiment 3, we provide evidence that explicit categorization of ambiguous items is a mechanism that leads to the single-category focus. Experiment 4 examines whether multiple category-use found in Experiment 1 arises from feature associations or category-level information.

Experiment 1

To examine whether subjects would integrate information across categories when predictions were not subject to explicit strategies, Experiment 1 used a cover task at test, the same-different task, in which predicting movement was incidental. Subjects learned four categories of moving shapes. Each category consisted of eight black moving geometric figures that varied in their shapes and movements. All shapes were presented in the center of a gray circle centered on a black computer screen. After a brief initial presentation, the shapes moved off the screen in a specific direction, which we will describe with clock directions. After category learning, subjects performed the same-different task in which they saw the same shapes they had learned, except the shapes now had diagonal stripes that were either tilted right or left (see Figure 2). The shapes moved along the same (uncertain) paths as during learning; however, their paths were now not entirely visible. After their initial presentation in the center, the shapes moved horizontally toward the edge of the computer screen (as in learning) but then disappeared behind an annulus, obscuring which direction the shape was going to move (see Figure 3). Shapes briefly reappeared from behind the annulus and then disappeared off the edge of the screen. When the shapes reappeared from behind the annulus, their stripes may have reversed their tilt (e.g., from left to right). The task was to report whether the tilt of the stripes was the same or different from when it appeared in the center of the screen.

In this test, subjects were never asked to predict direction or category, as they were only questioned about the stripes. However, since the shapes reappeared only briefly, looking close to where they reappeared—that is, making a prediction about the shape's direction—would aid performance. That is, knowledge of the shapes' categories gained during the learning task provided useful information about where the shapes would go, which, in turn, could help subjects make the same-different judgment. Indeed, the task was extremely difficult if one maintained central fixation. Position of eye gaze just before the shape's reappearance was the dependent measure, as a proxy for subjects' prediction of shape direction. However, since the task was to decide whether the object's internal feature had changed, and no question about direction was asked during either learning or test, we expected that people would not be making explicit predictions about direction, and so results should be different from cases in which they are asked to predict direction (Chen et al., 2014, Experiment 2).

There were two critical shapes of interest: squares and hearts. Each of these shapes belonged to one of two categories, the target or alternative categories. The *target category* is the category that the shape is most likely to be in. For example, there was a 67% chance that a square belonged to Category 1, the target category, and a 33% chance that it belonged to Category 2, the *alternative category* (i.e., there were eight squares in Category 1 and four in Category 2). In the target category, half of the squares moved in the 1 o'clock direction and half moved in the 5 o'clock direction. In the alternative category, the critical shapes moved in only one

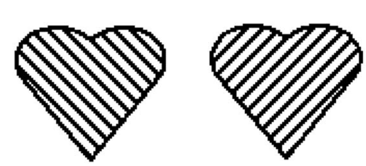

Figure 2. Example of stimuli used in the test phase of Experiments 1 and 2.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

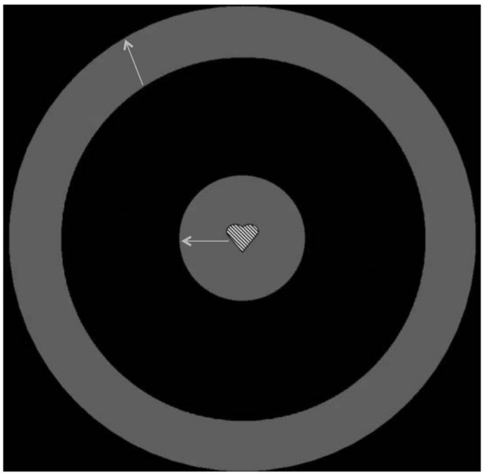

Figure 3. Illustration of the implicit induction task. The shape appeared in the center of the screen for 1 s. It then moved horizontally for .25 s and disappeared behind the annulus while traveling on its path (learned in Phase 1). Subjects reported whether the diagonal lines on the shape had changed when it reappeared. Arrows indicate the shape's visible path and did not appear in the experiment.

direction. In Condition 1, the squares moved to 1 o'clock; in Condition 2, which served to counterbalance the direction of the alternative category, they all moved to 5 o'clock. Therefore, if people only attend to the target category in predicting the direction of a new square, their average prediction should be around 3 o'clock. There are two (simple) ways in which an average of 3 o'clock could occur. First, subjects might consistently try to minimize the distance between the two possible locations of a square and look near 3 o'clock on most trials, or they could look equally often at 1 and 5 o'clock (probability matching).

If subjects attended to both the target and alternative categories, they should have a preference, because the alternative category (Category 2) would break the tie—in different directions in the two conditions. Thus, if people integrated information across categories they would, on average, shift their predictions depending on what condition they are in. That is, they should shift their prediction away from 3 o'clock toward the direction reinforced by the alternative category. This design was replicated for the heart stimulus using other directions (see Table 1). All subjects went through an identical learning phase in which they learned all four categories, based on the objects' shapes and direction of movement. Note that although our description focuses on the critical shapes, the categories included other shapes that did not have ambiguous categories (that will become relevant in Experiments 3 and 4), making the categories equal-sized.

Method

Participants. Subjects were 32 undergraduates at New York University who participated for course credit. Data from eight subjects were dropped for not providing recordable eye fixations before the shape's reappearance on at least five trials for both of the critical shapes. Subjects were randomly assigned to one of two between-subjects conditions, which served to counterbalance the direction of the alternative categories.

Materials. Stimuli for each category were eight black shapes approximately 1.75 to 2.5 cm in length, listed in Table 1. All exemplars of the same shape had identical dimensions. The same shapes were used during test except they had interior stripes oriented approximately 45° to the right or left (see Figure 2).

All stimuli were presented on the background of a light gray circle 30 cm in diameter centered on a black computer screen. Stimuli started in the center of the screen and then moved off the screen, disappearing once they moved beyond the border of the circle. Eye movements were monitored with the SR Research (Ontario, Canada) EyeLink 1000.

Procedure. The experiment consisted of three phases: (a) observation, (b) learning, and (c) test. A Macintosh computer presented the instructions and controlled all three phases. Eye movements were recorded during the test phase only.

Observation phase. Subjects were told that they would view four categories of moving shapes and were to learn what combination of shapes and directions belonged to each category for a memory test. During observation, all the shapes from each category (eight exemplars per category, see Table 1) were presented sequentially. Each shape appeared in the center of the screen for 1 s, then moved horizontally (toward 3 o'clock for shapes in Categories 1 and 2, toward 9 o'clock for Categories 3 and 4) for .4 s, and then moved toward its assigned clock direction for .95 s until it disappeared off the edge of the gray circle (see Table 1 for

Table 1
Category Structure Used in Experiments 1–3

	Category 1 (target for squares)		Category 2 (alternative for squares)		Category 3 (alternative for hearts)		Category 4 (target for hearts)	
Exemplar	Shape	Direction	Shape	Direction[a]	Shape	Direction[a]	Shape	Direction
1	Square	1	Square	1/5	Heart	7/11	Heart	7
2	Square	1	Square	1/5	Heart	7/11	Heart	7
3	Square	1	Square	1/5	Heart	7/11	Heart	7
4	Square	1	Square	1/5	Heart	7/11	Heart	7
5	Square	5	Rectangle	1/5	Diamond	7/11	Heart	11
6	Square	5	Rectangle	1/5	Diamond	7/11	Heart	11
7	Square	5	Rectangle	1/5	Diamond	7/11	Heart	11
8	Square	5	Rectangle	1/5	Diamond	7/11	Heart	11

Note. The direction entries are clock directions (1 = 1 o'clock, etc.).
[a] The first number refers to the direction in Condition 1, the second to Condition 2.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

directions). Each shape's category name ("Category 1," "Category 2," etc.) appeared in the center of the screen for the entire time it was on the screen. All exemplars from Category 1 were presented, then all exemplars from Category 2, and so on.

Learning phase. Subjects were next told that they would see the same items as in the observation phase. They were to classify each shape into one of the four categories by pressing a number key on the keyboard. At the beginning of each trial, a white fixation cross appeared in the center of the screen for 1 s. The shape then appeared in the center of the screen and moved as it did in the observation phase. There was no time limit on responding. After answering, the correct answer appeared for 1.25 s. After an error, subjects viewed a repeat display (without responding) of the moving shape with the correct category displayed. There were three learning blocks in which each of the 32 items (4 categories, each with 8 exemplars) was tested in random order. Because of the category uncertainty of the critical items (e.g., a square could be in Categories 1 or 2), subjects could get no more than 75% correct, assuming they chose the most likely category for all presented stimuli. In all experiments subjects had to reach at least 50% correct (chance = 25%) during the final block of learning to be included in analysis.

Test phase. The final phase of the experiment consisted of a 64-trial test (two blocks in which each of the 32 items was tested in random order) in which subjects had to perform the same-different task while their eye movements were tracked by the EyeLink 1000 using a 500 Hz sampling rate. Subjects were told that they would see the same items as in the previous phases except that the shapes would now move a little bit faster and have diagonal stripes on them. These shapes appeared in the center of the screen (for 1 s) and continued to move along the same path as in previous phases. However, there was now a black annulus on the screen such that the shape moved horizontally (for .25 s) and then disappeared behind the annulus for .7 s. The shape then reappeared from behind the annulus for .15 s before it disappeared from the screen. Recall that all stimuli were presented on a gray circle 30 cm in diameter. The annulus (24 cm in diameter) was centered on this image. Its center hole had a diameter of 8 cm (see Figure 3).

The stripes on a test object were either tilted left or right when the shape initially appeared in the center of the screen before moving behind the annulus (see Figure 2). The subjects' task was to report whether the direction of the stripes was the same or different when it reappeared from behind the annulus at the edge of the screen. The direction of stripes remained the same for half of the trials and changed for the other half. Subjects saw a 1.25 s feedback message. There were five practice trials before the test phase using a novel shape (a circle) to give subjects experience with the speed and task requirements. As shapes only briefly reappeared from behind the annulus, looking close to where shapes reappeared was beneficial. (Recall that horizontal movement for the critical shapes did not disambiguate its category, as the horizontal direction was the same for Categories 1 and 2, and Categories 3 and 4.)

Data analysis. The dependent measure was the fixation position recorded in the last sampling interval before a shape's reappearance. Responses for critical shape (square and heart) trials were coded such that a position exactly in between the two possible directions of the shape was 0 degrees, and a shift from that point toward the direction reinforced by the alternative category was coded as positive. For example, for the squares in Condition 1 (that might move to 1 o'clock or 5 o'clock), the 3 o'clock position was 0 degrees, the 1 o'clock position (the direction of the alternative category) was 60 degrees, and the 5 o'clock position was −60 degrees. In Condition 2, the latter values were reversed, as the alternative category reinforced the 5 o'clock direction in this condition. We obtained the mean fixation position for each subject by averaging the mean fixation position for squares and hearts. Thus, use of a single category (i.e., use of only the target category) is evidenced by an average prediction of 0 deg. Integration of information across categories is evidenced by a positive average prediction, as this represents a shift from 0 deg in the direction of the alternative category. (However, as will be seen, a positive or negative shift for any one condition could indicate a looking bias, so it is the overall looking shift that will represent multiple category use.)

Trials in which the fixation position was greater than 100 degrees or less than −100 degrees were not included in the analysis because the subject was fixated on the opposite side of the screen from where the shape traveled, indicating that the subject either forgot where the shapes went, or did not see the shape correctly before its movement. Additionally, trials where fixation was within the hole of the annulus were excluded from analysis. When subjects looked at the center of the screen while doing the task, they were effectively not making a prediction about direction. On average, fixation data for 18.9 trials (out of 64 total test trials) were not included in analysis for each subject because of the exclusion of trials based on the two criteria explained above and trials on which the eyetracker was not able to record fixation data (because of a subject blinking or moving, etc.).

Results

Subjects were on average 66.4% correct (chance = 25%) during their last training block, suggesting that they learned the categories quite well. (Recall that maximum performance was 75%, if subjects always classified ambiguous items into the most likely category.) Performance on the same-different task averaged 72%.

As explained above, integration of information across categories is evidenced by a shift from 0 deg in the direction of the alternative category, which we coded as positive. This is indeed what we found. The mean fixation position for the critical shapes, ($M = 7.5$ deg, $SD = 8.9$), was significantly greater than 0 deg, $t(23) = 4.1$, $p < .01$, $d = .84$, indicating that people's predictions of direction were integrated across the two categories. The mean fixation position was positive for 21 of the 24 subjects. (See Figure 4 for the distribution of eye positions on individual trials.) These results are consistent with those of Chen et al. (2014) and suggest that implicit induction promotes integration of information across categories.

There are at least two different response strategies that could lead to our finding that people used multiple categories and shifted toward the alternative category's direction when making implicit predictions. First, subjects' predictions could have been like a weighted mean. Eye position would have been between the two possible locations that the target shape might move to, but shifted toward the alternative category's direction. For example, subjects in Condition 1 may have generally fixated between 1 and 5 o'clock

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

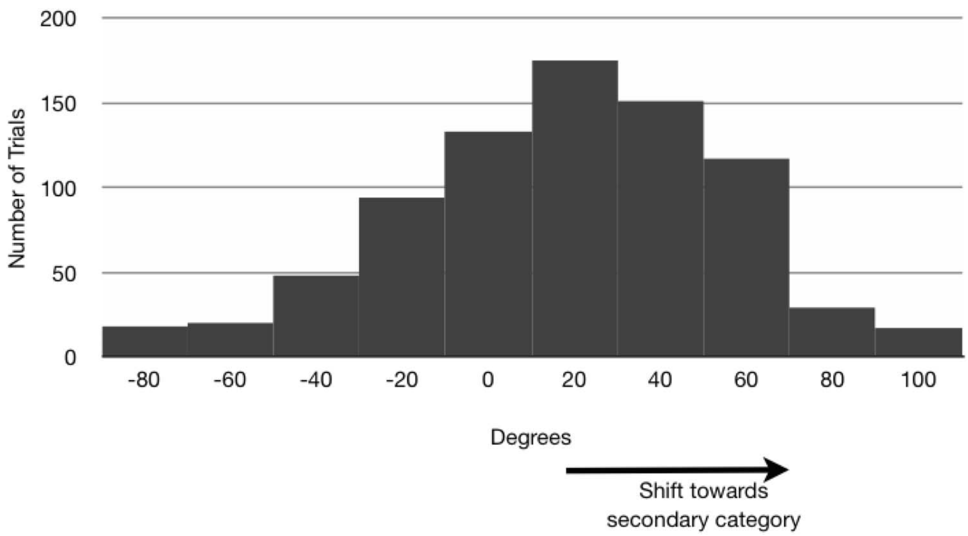

Figure 4. Histogram of eye fixation positions for individual trials in Experiment 1. Integration of information from multiple categories is evidenced by positive shifts from 0 deg. Each bin includes a 20 deg range and is labeled by the largest degree it includes.

on square trials, but closer to the alternative category's direction of 1 o'clock. Another strategy would be like probability matching. If subjects knew that the square could go to either 1 o'clock or 5 o'clock, they could have alternated between looking at these two locations roughly in proportion to how often the shape moved in that direction. That is, they would have fixated more often at the location of the alternative category (1 o'clock for square trial in Condition 1). The histogram of responses to individual trials (see Figure 4) is more consistent with the former explanation as the bulk of responses are in between 0 and 60 deg rather than a large cluster around 60 and a smaller cluster around −60 deg as a probability matching explanation would predict.

Perhaps subjects learned to change their eye movements only after practice in doing the task. To examine this possibility we compared the mean fixation position for the first and second blocks of testing. The positive shift in eye movements was significant in block 1, $t(23) = 3.3$, $p < .01$, $d = .68$, and in block 2, $t(23) = 3.6$, $p < .01$, $d = .76$, suggesting that subjects' use of multiple categories was not a result of learning during test. Additionally, the difference between the mean fixation positions for the first and second blocks was not significant (Ms = 6.2 and 8.8 deg, SDs = 9.1 and 11.9), $t(23) = 1.0$, $p > .05$, $d = .25$.

Noncritical shape analysis. There were two noncritical shapes, rectangles and diamonds, that did not enter into any of our hypotheses as their categorizations were certain (i.e., rectangles only appeared in Category 2 and diamonds in Category 3). We analyzed the results for these shapes to determine whether shifting eye fixation toward the direction of the shape's movement was in fact the strategy executed when there was certainty about the shape's trajectory. The true position that the noncritical shapes appeared at was 60 deg from the horizontal (0 deg). The average eye fixation position for these shapes was also greater than 0 deg, ($M = 13.7$ deg, $SD = 16.9$), $t(23) = 4.1$, $p < .001$, $d = .81$. Thus, subjects also shifted their eye fixation position toward the direction that the noncritical shapes moved.

Visual hemifield analysis. An unexpected result was that subjects tended to shift toward the direction of the alternative category more when the shape's alternative category reinforced a direction in the upper visual hemifield, *upward shapes* (squares in Condition 1 and hearts in Condition 2), than when the shape's alternative category reinforced a direction in the lower hemifield, *downward shapes* (hearts in Condition 1 and squares in Condition 2; see Figure 5). In fact, for downward shapes, subjects shifted slightly upward (reflected by a negative mean). A paired t test revealed that the average shift for upward shapes ($M = 23.3$ deg, $SD = 14.3$ deg) was significantly greater than for downward shapes ($M = -8.3$ deg, $SD = 18.4$ deg), $t(23) = 5.6$, $p < .01$, $d = 1.9$.

We speculate that this tendency to shift fixation more in the direction of the alternative category when its associated direction is in the upper hemifield is because visual acuity and performance on a variety of tasks is better in the lower than the upper visual hemifield (Edgar & Smith, 1990; Levine & McAnany, 2005). Thus, subjects would not have to fixate as close to the shape's reappearance position to do well on the task when the shape reappeared in the lower hemifield as when the shape reappeared in the upper hemifield. Indeed, performance on the same-different task was equivalent for the upward and downward shapes even though eye fixations were further away in the downward condition. Furthermore, the amount of shift toward the alternative category's direction on a given trial was positively correlated with performance on that trial for upward shapes, $r = .1$, $p < .05$, but not correlated with performance on downward shape trials, $r = -.06$, $p > .05$. This suggests that shifting fixation toward the alternative category's direction improved performance when shapes reappeared in the upper hemifield, but not when they reappeared in the lower hemifield. We are not suggesting that subjects shift toward the upper hemifield because they explicitly know that acuity is worse there or that looking in that direction will improve performance. Rather, it seems likely that the visual system is well-trained

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

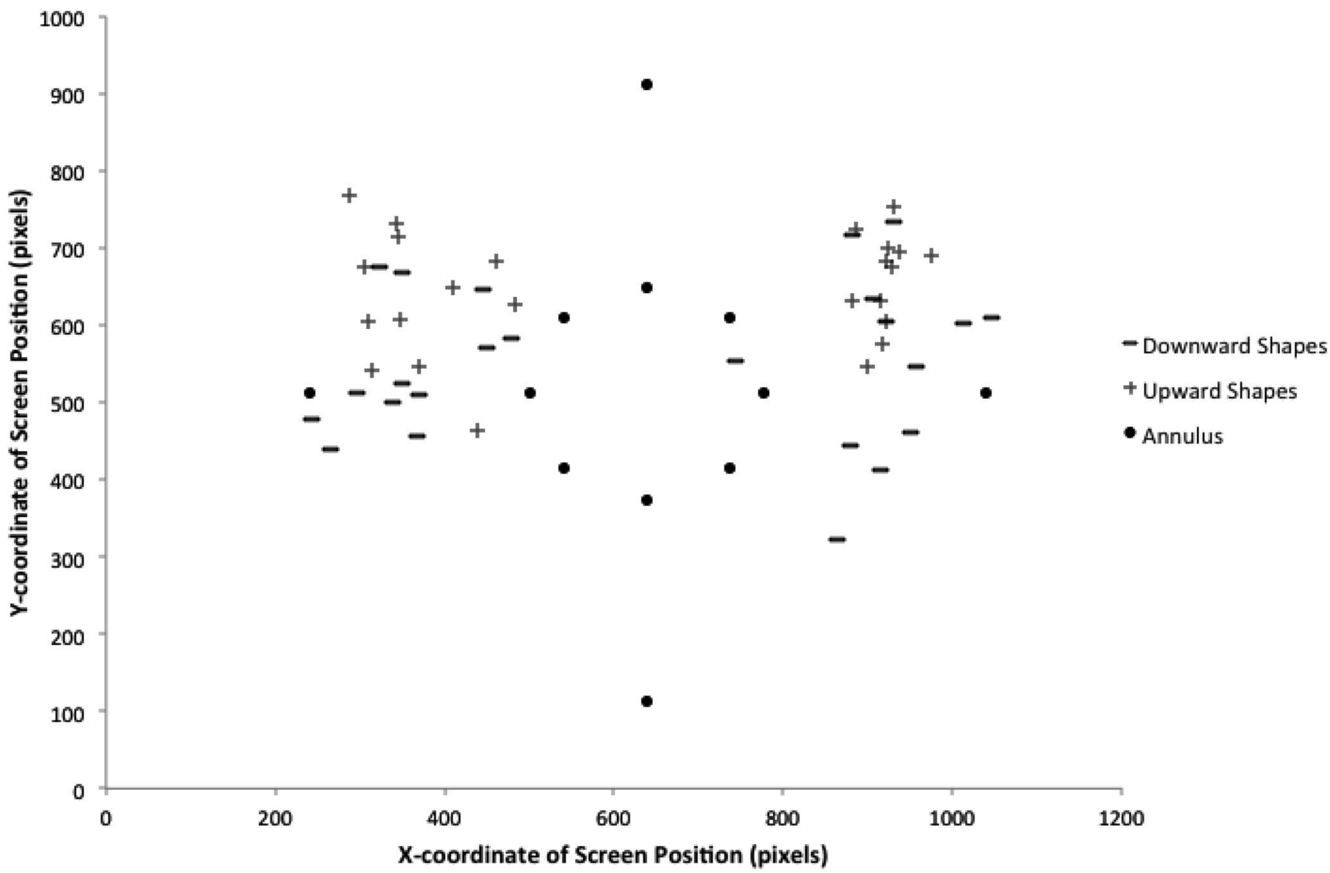

Figure 5. Average fixation position for each subject for downward shapes (alternative category reinforces direction in lower hemifield, minus signs) and upward shapes (alternative category reinforces direction in upper hemifield, plus signs). Although there is a general bias to look upwards (the midpoint is Y = 512), the plus signs are generally higher than the minus signs.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

to fixate in areas in which acuity is poor and that this tendency is automatic.[1] In visual search tasks, subjects have been found to fixate more in areas where visual fields are less sensitive despite the fact that the target is not more likely to appear in these areas (Najemnik & Geisler, 2005, 2008).

Discussion

The results of Experiment 1 indicate that subjects integrated information across categories when making their predictions implicitly. In the same-different task, the prediction was only incidental to the task—direction was never mentioned to the subjects; their task was only to report about a feature completely independent of direction, the tilt of the stripes. Thus, the predictions of direction we measured via eye fixation position were likely outside of strategic control. Experiment 2 examines this issue further.

Subjects' bias to look up rather than down slightly complicates the results. There was a reliable effect of looking in the direction indicated by the alternative category when averaged across hemifields, indicating category integration. One way to think about the results shown in Figure 5 is that the + signs, indicating upward shapes, are generally higher than the − signs, indicating downward shapes. This shows that the alternative category had an effect above and beyond the overall upward bias.

The category structure of Experiment 1 was such that the target category did not give a clear answer for the direction prediction. That is, within the target category a shape was equally likely to go in two directions. This category structure likely does encourage use of the alternative category as a tie-break of sorts compared with structures where the target category gives a clear prediction (Murphy & Ross, 2010b). However, the category structure alone does not account for the results of Experiment 1 (Chen et al., 2014,

[1] Another factor that may account for the difference in shift between the two visual hemifields is that in Western cultures there is a habit to scan from top to bottom (Wickens & Hollands, 2000). Thus, when the alternative category reinforced a direction in the upper hemifield, the corresponding shift was congruent with that habit. When the alternative category reinforced a direction in the lower hemifield, the corresponding shift was incongruent with that habit and perhaps, in part, led to the lack of shift found in the downward direction. While this top-to-bottom scanning habit may have contributed to the shift asymmetry found between the visual hemifields, it would not explain why amount of shift correlates with performance on the same-different task in the upper but not lower hemifield.

Experiments 2 and 3). Experiment 2 investigates whether explicit decisions would produce the same results.

Experiment 2

The finding that subjects integrated information across categories in Experiment 1 supports the conclusion that subjects may use category information more normatively when induction is implicit because these predictions are outside of strategic control. What would happen if people made their predictions explicitly, simply stating where they thought shapes would go? Chen et al. (2014) used the identical category structure and learning procedure and tested explicit predictions of the shape's direction. In three explicit conditions tested, they not only did not find correct Bayesian integration of information across categories, they actually found that people's predictions were in the *opposite* direction from what integration of categories would predict.[2] There would be little point in replicating these verbal predictions here, which would be the fourth experiment to test this category structure with verbal predictions. However, we did explore whether the present task encouraged people to adopt an explicit eye movement strategy that would not have been present in the earlier studies.

One possibility is that the integration of information found in Experiment 1 was part of a strategy developed during practice. Because subjects performed five practice trials before the test trials, they may have realized that they would perform better when they looked closer to the direction reinforced by the alternative category. Perhaps they were able to develop a strategy to look in the direction of the most frequently observed direction (though why such a strategy should include a difference between lower and upper visual fields is very unclear, since most people are not aware of this asymmetry). That is, perhaps subjects' strategy reflected an explicit prediction of direction. Alternatively, the learning process itself may have led subjects to look up for some shapes and down for other shapes, which knowledge they then used in the test phase. As we have proposed that implicit induction leads to integration of information by avoiding overt strategies that are often used in traditional category-based induction tasks, we must determine whether subjects in Experiment 1 were aware of the strategy used to perform the same-different task. According to our proposal, subjects should not be able to report an integration strategy if asked how they would perform the same-different task from Experiment 1. Experiment 2 examines this issue.

In Experiment 2, subjects completed the full learning procedure of Experiment 1. They then saw the same example trials of the same-different task used in Experiment 1. At test, they saw the initial figure on each trial but did not actually perform the same-different task. Instead, they moved their eyes to where they believed they would look to best perform the task. This question sampled subjects' explicit beliefs about where they would look. If the results match those of Experiment 1, this would suggest that the fixations were the result of an explicit strategy.

Method

Participants. Subjects were 21 New York University undergraduates who participated for course credit. Data from four subjects were dropped for not providing recordable eye fixations on at least three trials. One more subject was dropped for not reaching the performance criterion during learning.

Materials and design. The materials and design were identical to those of Experiment 1.

Procedure. The observation and learning phases were identical to those of Experiment 1. As in Experiment 1, eye movements were only recorded during the test phase. The test phase consisted of a 16-trial test in which subjects reported where they would look to best do the same-different task that subjects in Experiment 1 performed. The number of trials was smaller than in Experiment 1, because the same four questions repeated with no feedback (e.g., subjects were asked where they would look to best perform the same-different task for a square multiple times, same for hearts, rectangles, and diamonds). Thus, we did not want to indicate to subjects that we were looking for a different answer by repeatedly asking them the same question. In Experiment 1's test, the critical shapes went in different directions and differed in whether the shading changed, so those trials did not appear to be exactly the same.

Subjects saw the same five practice trials used in Experiment 1 and then were told that they would not be doing the task but rather reporting where they would look just before the shape's reappearance from behind the annulus to best do the task. To keep the dependent measures of the two experiments similar, we used eye fixation to indicate this prediction. A white dot on the display indicated where the subjects were looking. Their task was to look at the location on the screen that they thought would be best to do the same-different task they had just observed. At practice, they then saw a test screen (gray circle with the annulus) and were instructed to look around the screen to get a sense of how the white dot corresponded to their eye gaze.

The test phase consisted of four blocks in which each of the four shapes (square, heart, rectangle, and diamond) was tested once in random order (except that shapes were not queried in two consecutive trials). Each test trial started with the presentation of the shape in the center of the screen for 1 s. It then moved horizontally for .25 s until it disappeared behind the annulus (the shape never reappeared). Subjects then saw the white dot that marked their eye gaze on the screen. To report where they thought they would look while doing the task, subjects moved their eyes until they were satisfied with the location of the white dot and then pressed the enter key. The white dot stayed on the screen for 1.25 s so that the subjects could see their answer.

Results

Subjects were on average 68.2% correct (chance = 25%) during their last training block, near the 75% maximum and virtually identical to the learning rate in Experiment 1.

As in the analysis of Experiment 1, subjects' responses for the critical shapes were coded such that the time corresponding to the point exactly in between the two possible directions of the shape was 0 degrees (3 o'clock for squares and 9 o'clock for hearts, and a shift toward the direction reinforced by the alternative category

[2] The exact effect depended on whether all trials were included or only those in which subjects selected the target (most likely) category. When all trials were included, their predictions were slightly though not significantly negative (−4.6, 0.7, and −5.2 deg). When subjects predicted the object was in the target category, their predictions were strongly counter to the Bayesian prediction (−32.5, −15.8, and −27.1 deg). See Chen et al. (2014) for detailed discussion.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

was positive). Again, positive scores would indicate use of multiple categories. To find the mean prediction (the amount of shift from 0 deg toward the alternative category) for each subject, we again calculated the mean prediction for each critical shape and took the average of the two.

Subjects' responses showed no evidence of integration of information across categories and did not match the results from Experiment 1. The mean prediction (M = 0.2 deg, SD = 2.9 deg) was not significantly different from the average observed direction for the shapes in their target category only (0 deg), $t(15) = 0.2$, $p >$ .05, d = .07, suggesting that subjects did not base their responses on multiple categories. This is not merely an effect that did not differ from 0: Subjects chose locations around 0 deg the vast majority of the time. In fact, 84% of all responses were within 10 deg of 0 deg. In Experiment 1, only 25% were in this range. Subjects clearly did not know that they would look at the locations where the shapes were likely to reappear. In short, most subjects thought they would not move their eyes much to perform the task, contrary to the actual fixations of Experiment 1.

Results for the noncritical shapes revealed that responses for these shapes did differ significantly from 0 deg, albeit it by a very small amount (M = 1.5 deg), $t(15) = 2.6$, $p < .05$. (Compare this to an actual shift of almost 14 deg in Experiment 1.) Thus, for the noncritical shapes, subjects had some ability to report that they would shift away from 0 deg toward the direction of the direction of the alternative category. However, unlike the critical shapes, a shift toward the direction of the alternative category does not involve integration of information from multiple categories as noncritical shapes appeared in only one category, the alternative category.

Discussion

The results of Experiment 2 further suggest that the integration of information across categories found in Experiment 1 was not the result of an explicit strategy to look where the object was likely to go. Subjects were unable to report the strategy that they would use to perform the same-different task. Indeed, their fixations bore no resemblance to where subjects looked while actually performing the task in Experiment 1.

The results of Experiment 2 build on previous results from experiments that used the same category structure. Those results showed that verbal prediction of direction also did not show the integration of information across categories (Chen et al., 2014). Thus, it seems unlikely that the integration of information across categories found in implicit induction were a result of strategic decisions about where to look, as neither explicit predictions of direction nor explicit reports of prediction task strategy (Experiment 2) showed any evidence of this. Additionally, as the same category structure and learning procedure was used in Experiments 1 and 2, these results further suggest that the integration of information found in Experiment 1 was not simply a result of the category structure or learning procedures used.

Subjects in Experiment 2 performed the same-different task with novel stimuli to illustrate it but did not do the task with the critical stimuli. One possibility (suggested by a reviewer) is that subjects in Experiment 1 developed an explicit strategy to look in the most likely direction as a result of experience in doing the task. That is possible, but it does not seem consistent with the lack of a practice effect in Experiment 1's test—the shift toward the most likely direction was no larger for the first block of test trials than the second. Subjects would have had to realize almost immediately, "I'd better look where the object is going to go" and then formed the explicit goal of using both categories. Additionally, previous research has found that subjects' verbal predictions of where objects will go also do not show evidence of multiple-category integration (as reviewed above; Chen et al., 2014). There is no obvious reason why people's explicit predictions should be different in these two cases.

To be clear, we cannot draw a strong conclusion about just what subjects in this experiment did think about where the objects were going to go. The conclusion is merely that they did not form an explicit eye-movement strategy for how to do the task that reflected anything like what their actual performance was.

Experiment 3

The distinction between implicit and explicit cognition is a controversial one, especially regarding whether these are two distinct systems, with distinct processes and neural bases (e.g., Evans & Stanovich, 2013; Kruglanski & Gigerenzer, 2011). We are not committed to this distinct systems claim. Our view is that the distinction is a useful heuristic that has led to many interesting findings. Explicit and implicit processes are probably family resemblance categories, each sharing certain properties that sometimes overlap with processes in the other category. Therefore, we ask what typical properties of the (proposed) implicit tasks we have used that lead to use of multiple categories and why explicit tasks generally lead to single-category use. It may also be possible to find a seemingly implicit task that nonetheless allows people to focus on a single category and a seemingly explicit task that encourages multiple-category use, if it has the correct psychological properties.

Therefore, we have considered what aspects of the tasks we have tested lead to the differences we have observed. A clear difference between our implicit and explicit tasks is that explicit prediction tasks either elicit an explicit categorization or allow time for one before the induction. Thus, subjects may commit to this initial categorization even though it is uncertain. As our proposal is that implicit induction prevents the use of strategies that lead to the systematic disregard of information from less likely, alternative categories, initial categorization is a likely part of this filtering process. Implicit predictions may avoid a single-category focus by not encouraging this initial categorization and/or by operating under time pressure, making it unlikely that subjects categorized spontaneously. Chen et al. (2014, Experiment 1) found some suggestive evidence that faster responding was associated with greater use of multiple categories. Thus, the lack of initial categorization may allow implicit inductions to avoid the single-category focus found in explicit induction. In explicit tasks, requiring subjects to process all of the presented categories also reduces the single-category focus (Murphy & Ross, 2010b).

Experiment 3 was designed to test the impact of this variable on how category information is used. In Experiment 3, subjects learned the same categories with the same procedure as in Experiments 1 and 2. At test, they again performed the same-different task from Experiment 1, except that on each trial they first viewed the (static) shape in the center of the screen and were asked what

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

category the shape was most likely to be in. Once subjects responded (they received no feedback on their categorization), the shapes moved rapidly off the screen as they did in Experiment 1, and subjects performed the same-different task. Eye fixation position just before the shapes' reappearance from the annulus was again used as the dependent measure. Thus, this task was identical to the implicit induction task from Experiment 1 except for an initial categorization before the induction task. Subjects' predictions (eye movements) were still implicit in that we never asked what direction the object was going to move, and the overt task was the same-different judgment. (The added categorization task also does not test direction of movement.) If categorization promotes the disregard of information from less likely alternatives, subjects' predictions in this task should show less, or no, evidence of integration of information across categories.

Method

Participants. Subjects were 31 New York University undergraduates who participated for course credit. Data from seven subjects were dropped for not providing recordable eye fixations on at least five trials for both of the critical shapes. Three more were dropped for not categorizing the shapes in their target categories during the test phase.

Materials and design. The materials and design were identical to those of Experiment 1.

Procedure. The procedures of the observation and learning phases were identical to those of Experiment 1. Again, eye movements were only recorded during the test phase.

The test phase consisted of 64 trials in which subjects performed the same-different task used in Experiments 1 and 2 except that they categorized the shape before each trial. Each trial started with a shape presented in the center of the screen and the question, "What category do you think the shape below most likely belongs to?" Subjects responded by pressing the number key on the keyboard that corresponded to the category. There was no time limit for responding. Once subjects responded, the shapes moved along the same path as in previous phases with the same timing used in Experiment 1 (horizontally for .25 s, disappears behind annulus for .7 s, reappears from behind annulus for .15 s before disappearing from the screen), and subjects reported if the tilt of the stripes was the same or different than the initial presentation. Subjects saw a 1.25 s feedback message. If subjects responded to the categorization question in less than 1 s, then the shape and prompt stayed on the screen for a total of 1 s before the shape started to move. This is because in Experiment 1 the shapes appeared in the center of the screen for 1 s before movement. Subjects saw the same five practice trials used in Experiment 1 except that, as with the test trials, the categorization question was asked before the same-different task. Subjects were told that since the shapes in the practice trials (circles) were not in any of the learned categories, they should categorize them into Category 5.

Results

Subjects were on average 68.1% correct (chance = 25%) during their last training block, near the 75% maximum, suggesting that they learned the categories well. During test, they categorized the critical shapes into their target category 60.4% of the time and into their alternative category 37.0% of the time. (Recall that critical shapes appeared in their target 67% of the time and in their alternative category 33% of the time, so this approaches probability matching.) Subjects categorized these shapes into categories that they did not belong to only 3.6% of the time. Performance on the same-different task averaged 69.4%. This is similar to performance without the initial categorization (72%) in Experiment 1, so the categorization question did not unduly influence the same-different task.

We examined subjects' responses for the critical shapes depending on how the shape was categorized before performing the same-different task. If explicit categorization influenced inductions, predictions should be different when subjects classified into the target category than when they classified into the alternative category. Specifically, when the shape was classified into the target category, we would expect predictions not to be significantly different from 0 deg. When the alternative category was picked, we would expect prediction to be significantly different than 0 deg. This is indeed what we found. The mean prediction when the critical shape was classified into its target category was close to 0 deg (M = 1.0 deg, SD = 15.8 deg), $t(20) = 0.3$, $p > .05$, $d = .06$, suggesting that these predictions were only based on the target category. The mean prediction when the critical shape was classified into its alternative category (M = 13.0 deg, SD = 15.8 deg) was significantly greater than 0 deg, $t(16) = 3.4$, $p < .01$, $d = .82$. (There are fewer degrees of freedom for this analysis because not all subjects classified the critical shapes into their alternative category.) A paired t test revealed that the difference between the mean fixation positions for target and alternative categorizations was marginally significant (Ms = −1.6 and 8.0 deg, SDs = 16.0 and 17.9), $t(16) = 1.7$, $p = .1$, $d = .57$ (only subjects who had both target and alternative categorizations were included in this analysis).

Thus, when subjects classified shapes into their target category, their induction appears to be based on only the target category—their prediction of direction was not significantly different from the average of the two directions associated with that category (0 deg). To examine if predictions on trials in which shapes were classified into their alternative category were also based on only a single category (i.e., based only on the alternative category), we compared these predictions to direction predictions of noncritical shapes. Because noncritical shapes appeared in only one category (the critical shapes' alternative categories—see Table 1), these predictions should be based on only the alternative category. If there is no difference between the predictions, this would suggest that subjects based their predictions on only the alternative category when they categorized the critical shapes in this category. This is indeed what we found. The mean prediction for noncritical shapes (M = 19.5 deg, SD = 25.0 deg) was not significantly different from the mean prediction for critical shapes when they were classified into their alternative category (M = 13.0 deg, SD = 15.8 deg), $t(16) = .90$, $p > .05$, $d = .32$.

Discussion

The results of Experiment 3 suggest that explicit categorization is a critical difference that leads to the different use of categories in implicit and explicit induction. When prompted to categorize the shapes before doing the implicit induction task, subjects no longer

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

showed reliable integration of information across categories found in Experiment 1. Subjects instead based their predictions on only a single category. When subjects categorized the shape into its target, they based their prediction on only that category; when they categorized the shape into its alternative, they based their prediction on only the alternative category.

It is worth emphasizing that making a prediction consistent with only the initial categorization is nonnormative. Subjects clearly knew that the categorization was uncertain (from learning and from the fact that they categorized the same shapes into different categories during test). They received no feedback on their categorization, so the shape's categorization remained as ambiguous after their categorization as it was before it. Suboptimal predictions because of overcommitment to an uncertain initial categorization has been found in previous research on category-based induction using verbal predictions (Chen et al., 2014; Lagnado & Shanks, 2003; Murphy et al., 2012). The present result is unique in that the explicit categorization influences *implicit* predictions (eye movements), suggesting that these initial categorizations have a broad influence on our subsequent judgments.

How exactly did categorization influence eye movements in this task? There are two main possibilities, one more implicit and the other more explicit. The first claims that initial categorization would activate the locations associated with that category. As a result, when the object disappears, people are more likely to look in those directions, even though their initial categorization is uncertain. Here, there's no need for an explicit prediction of direction; subjects' predictions of direction are incidental to the same-different task and their implicit predictions are affected by the activation from the initial categorization.

The second possibility is that once people consciously selected one of the categories, they are more likely to make an explicit prediction of direction. As a result, this single-category focus is actually a reappearance of explicit effects found in past research with paper-and-pencil responses. Either of these possibilities is consistent with our claim that initial categorization is important in determining the use of multiple categories. The possibility referring to an explicit prediction, however, is not entirely consistent with the results of Experiment 2 in which we asked people where they would look to perform the task. In Experiment 2, people generally kept their eyes at the midline, with little variation (SD = 2.9 deg), whereas when actually doing the task in Experiment 3, people moved them much more (before the object's reappearance; SD = 15.8 deg). Also, for the noncritical shapes, subjects in the Experiment 3 made large shifts (M = 19.5 deg), whereas in Experiment 2, they again hardly shifted at all (M = 1.5 deg), even though there was no uncertainty about where the object would go. It is possible that the conscious strategy obtained in the present experiment was different from that of Experiment 2, but the eyetracking performance observed here seems overall more similar to that of the implicit eyetracking (except for the single-category aspect) found in Experiment 1.

We are suggesting that the focus on single categories found in past research is also due, in part, to initial categorization. In Experiment 3, this initial categorization took the form of an explicit categorization question and a verbal response. However, it should be understood that in some induction situations categorization will be spontaneous. For example, in paradigms that use familiar categories that are mentioned in a scenario, it seems likely that people will use those categories even if not forced to do so by a question. For example, if a scenario describes a person as probably the real estate agent you are expecting but possibly a cable TV repair person, people seem to choose one or the other category even if the categorization question appears after the predictions (e.g., Malt et al., 1995). Similarly, when presented with visual displays in which categories are separated into distinct groupings, people very likely decide which category is correct even if they are not immediately asked about it (Murphy & Ross, 1994, Experiments 2 and 5; Murphy & Ross, 2005, Experiment 4). These paradigms have in common two important factors. First, the categories are salient, either by being entrenched through familiarity or by being perceptually separated. Second, they are untimed tasks in which people can think about the categories as much as they would like.

In our original implicit induction task (Chen et al., 2014), the task moves very quickly, and people are never asked (before or after) about the items' category. Not only is classifying the object not required, doing so quickly would probably be quite difficult while still catching the object—or while doing the same-different task in Experiment 1. Thus, it seems doubtful that subjects engaged in spontaneous classification. However, this discussion does point out that whether people have classified items is not always straightforward to determine. In the present experiments, it seems very likely that people were not classifying in Experiment 1 but were in Experiment 3; in other situations, investigation may be necessary to ascertain whether people classified before making a prediction.

Experiment 4

The multiple-category use found in Experiment 1 is also consistent with a feature-level strategy (e.g., determining what direction a new square will travel in by only considering square exemplars) rather than a category-level strategy like that described in the Introduction (see Griffiths et al., 2011; and Newell et al., 2010, for similar ideas). As it has been suggested that people are more likely to use multiple categories in induction when they can use a feature-conjunction strategy (Papadopoulos et al., 2011), it may be that the multiple-category use found in Experiment 1 was based on a feature-feature associations (e.g., squares tend to go the lower right) rather than category-level information (e.g., objects in Category 2 tend to go to the lower right). Thus, it may be that subjects were not using multiple categories (or even categories at all), but instead were only using features to perform the implicit induction task. This account is consistent with our proposal that "people use multiple categories" in implicit induction, in the sense that information from multiple categories influence performance, even if people are not actually considering categories—for example, if subjects were only considering squares, given the positive shift from zero degrees found in Experiment 1, they must have been considering squares from *both* Categories 1 and 2. However, it is obviously of interest to know whether people are actually using category-level information or only feature information (see Murphy & Ross, 2010a; Newell et al., 2010). Indeed, one could argue (as a reviewer suggested) that the experimental categories were too weak to serve as actual categories with one observation block plus three learning blocks and that in fact people categorized the items by their shape instead. If so, then the results of Experiment 1 arise from a single-category prediction, shape to direction, contrary to our conclusion.

Griffiths, Hayes, and Newell (2012) examined a number of conditions of learning and responding to categories in a similar explicit task (predicting colors from shape). They concluded that when categories are learned, people generally use category-level information. When people do not respond based on learned categories but instead on the basis of a visual display of the category information, they tend to use feature-level information (though a significant number continue to limit their responses to a single category). Their study would suggest that our subjects were responding based on category information, given that induction was based on memory. However, that conclusion still assumes that the learned categories (i.e., Categories 1–4) were the ones subjects used. If instead people were treating shape as the category, then our experiment would not have provided a test of whether multiple or single categories were used.

To explore which strategy subjects use in this task, we constructed a new category structure in which the category-level and feature-based strategies made different predictions of eye fixation positions. We accomplished this by changing the direction of motion of the noncritical shapes, which changed the category-level direction but not the direction associated with the given features (the critical shapes). As in the category structure used for the previous experiments, square and hearts are the critical shapes. The target category for both shapes remained the same (e.g., in Category 1, half of the squares moved in the 1 o'clock direction and half moved in the 5 o'clock direction). In the alternative category, the critical shapes moved in only one direction (as in the previous category structure). In Condition 1, all the squares moved to 1 o'clock; in Condition 2, which served to counterbalance the direction of the alternative category, they moved to 5 o'clock. Thus, if subjects base their predictions on only squares (a feature-based strategy), a positive shift from 0 degrees would result.

Unlike the previous category structure, the noncritical shapes (rectangles and diamonds) moved in the opposite direction of the critical shape in the alternative category (in the previous structure critical and noncritical shapes moved in the same direction). In Condition 1, the rectangles moved to 5 o'clock; in Condition 2, they moved to 1 o'clock (see Table 2 for full category structure). Because there are an equal number of critical and noncritical shapes in the alternative category, if subjects use a category-level strategy, predictions should be around 0 degrees, in contrast to the results of Experiment 1 and the implicit prediction results of Chen et al. (2014).

This design leads to distinctive predictions of the two hypotheses (which can be seen in a comparison of Tables 1 and 2). If the results of Experiment 1 were because of feature associations, then there should be a positive deviation from 0 degrees in eye fixations in Experiment 4, and there should not be a significant difference between these results and those of Experiment 1 (that had identical feature-direction pairings for the critical shapes). If the results of Experiment 1 were because of category-level effects, then the present eye fixations should not be significantly different from 0 degrees, and the mean fixations should be significantly different from those of Experiment 1.

Method

Subjects were 35 New York University undergraduates who participated for course credit. Data from six subjects were dropped for not providing recordable eye fixations on at least five trials for both of the critical shapes. One more was dropped because of technical error. The materials and procedure were identical to those of Experiment 1 with the exception of the modified category structure (see Table 2).

Results

Subjects were on average 64.1% correct (chance = 25%) during their last training block, nearly identical to the 66% correct of Experiment 1. Performance on the same-different task averaged 72%, the same as Experiment 1. Thus, the change in category structure did not alter learning or overall performance on the test.

Subjects' responses for the critical shapes were again coded such that the time corresponding to the point exactly in between the two possible directions of the shape was 0 degrees (3 o'clock for squares and 9 o'clock for hearts). A shift toward the direction reinforced by the critical shape in the alternative category was positive. Thus, an average prediction (fixation placement) greater than 0 degrees would indicate that subjects used a feature-level strategy. Scores of 0 would indicate the use of a category-level strategy.

The mean prediction for critical shapes was not significantly different from 0 deg ($M = -.89$ deg, $SD = 8.0$ deg), $t(27) = 0.6$,

Table 2
Category Structure Used in Experiment 4

	Category 1 (target for squares)		Category 2 (alternative for squares)		Category 3 (alternative for hearts)		Category 4 (target for hearts)	
Exemplar	Shape	Direction	Shape	Direction[a]	Shape	Direction[a]	Shape	Direction
1	Square	1	Square	1/5	Heart	7/11	Heart	7
2	Square	1	Square	1/5	Heart	7/11	Heart	7
3	Square	1	Square	1/5	Heart	7/11	Heart	7
4	Square	1	Square	1/5	Heart	7/11	Heart	7
5	Square	5	Rectangle	5/1	Diamond	11/7	Heart	11
6	Square	5	Rectangle	5/1	Diamond	11/7	Heart	11
7	Square	5	Rectangle	5/1	Diamond	11/7	Heart	11
8	Square	5	Rectangle	5/1	Diamond	11/7	Heart	11

Note. The direction entries are clock directions (1 = 1 o'clock, etc.).
[a] The first number refers to the direction in Condition 1, the second to Condition 2.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

$p > .05$, $d = -.11$. Thus, there is no evidence that subjects were using a feature-level strategy when making these predictions. As discussed above, the feature-level strategy and the category-level strategy make different predictions about whether there will be a difference between the average fixation position in Experiments 1 and 4. A feature-level strategy would predict a positive shift from zero degrees in both experiments; a category-level strategy would lead to a positive shift in Experiment 1 but no shift in Experiment 4. The latter is indeed what we found: The average shifts in Experiment 1 ($M = 7.5$ deg) and Experiment 4 ($M = -.89$ deg) were significantly different, $t(51) = 3.6$, $p < .001$, providing evidence that directions were computed at the category-level.

Surprisingly, the mean prediction for the noncritical shapes was significantly different from 0 ($M = 8.4$ deg, $SD = 14.9$ deg). The category-level strategy would also lead to an average prediction of 0 degrees for the noncritical shapes as the half the exemplars in these categories (Category 2 and 3) moved toward 60 deg and half moved toward -60 deg.[3] Thus, unlike the critical shapes, there is evidence that subjects used a feature-based strategy for the noncritical shapes.

Discussion

The results of Experiment 4 indicate that subjects were not using a feature-level strategy when making category-based inductions under uncertain categorization in our task. Subjects did not shift their predictions for the critical shapes away from 0 degrees even though the feature-direction associations would have predicted that, as the same associations (for critical shapes) were present here as in Experiment 1. The comparison between the results of Experiments 1 and 4 provide some evidence that subjects were using a category-level strategy. The prediction of about 0 degrees found in Experiment 4 is, of course, also consistent with a single-category strategy (assuming items are always classified into their target category), but the results of Experiment 1 and previous studies which have shown clear evidence of multiple-category use in this task and another similar task (Chen et al., 2014) suggest that this interpretation is unlikely. It is very unclear why a few changes in noncritical stimuli would lead to such a large change in strategy. Overall, the results suggest that implicit category-based induction under uncertain categorization is, in fact, category-based induction (rather than feature-based induction), at least, under the conditions studied here. They also confirm that subjects were not treating the critical shapes as the category, or else the results of Experiment 1 would have been replicated.

We would not necessarily extend these results to address the issue of category and feature use in different paradigms, as studied by Griffiths et al. (2011), Murphy and Ross (2010a), or Papadopoulos et al. (2011), for example. In those experiments, whole categories were visually presented, and people could easily focus on stimuli that had the given feature. For example, when predicting the color of a new square, subjects could examine each presented square and ignore the other shapes. In the present paradigm, where category knowledge was learned and then tested from memory, such a strategy would be much more difficult, as individual exemplars would have had to be memorized and then quickly retrieved. Griffiths et al. (2012) drew the same conclusion after comparing judgments from memory versus visual display. When making predictions from memory, it is probably more efficient to use representations of whole categories, which is of course one of the reasons often given for why we have categories (Rosch, 1973).

The results for the noncritical shapes suggest that, for these shapes, feature associations were controlling eye movements. In retrospect, this may not be surprising, because every noncritical shape always goes in the same direction throughout learning and test. That is, every rectangle might go to 1 o'clock. Thus, unlike the cases we have been focusing on in which categorization and features are probabilistic, here the feature associations are much stronger and more reliable. As Griffiths et al. (2012) document, variations in learning procedure and category structure influence whether category-level and feature-level information control prediction. Presumably the present difference between critical and noncritical shapes represents another variation, in which feature-level information dominates category-level knowledge when features are perfectly associated. Of course, this conclusion must be tentative, as the experiment was not designed to test that effect, and there are differences between the critical and noncritical items in our design (frequency being the most notable).

General Discussion

The experiments in this article examined implicit induction and why it is able to avoid reasoning biases often found in explicit induction. In comparing different ways of making inductions we observed four main results. First, we found that when making implicit predictions under uncertainty people integrated information across categories (Experiment 1). Second, we found evidence that this integration was not a result of strategic decisions. Subjects' explicit reports of their strategy for performing the implicit induction task showed no evidence of integration of information across categories (Experiment 2). Third, we found evidence that implicit induction may allow for more normative integration of category information by avoiding explicit categorization before induction (Experiment 3), which may have the effect of focusing attention on the selected category. Finally, we also discovered that in this paradigm, subjects were using category-level information rather than simply activating feature-feature associations (Experiment 4).

We have used the terms *implicit* and *explicit* because we believe they efficiently communicate the important differences between the conditions: fast decisions based on associations with little strategic input versus slower decisions mostly using consciously selected strategies. This distinction is obviously reminiscent of Sloman's (1996) important distinction between implicit and explicit reasoning, as well as various dual system views of reasoning (Evans, 2007; Kahneman, 2011). However, we do not know whether our two conditions draw on two distinct cognitive systems, and the results of Experiment 3 may resist a simple explanation of this sort. There are many different cognitive processes that could carry out predictions of different kinds, and each might have its own empirical profile. Clearly, every such difference does not imply distinct cognitive systems. Our distinction does not critically rely on a dual systems approach or a claim that the implicit predictions made in this article are completely outside of

[3] A comparison with the results of Experiment 1 is not informative in this case, because the category-level strategy and the feature-level strategy in Experiment 1 suggest the same prediction for the noncritical shapes.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

consciousness, but it shows that there are different reasoning processes that underlie different tasks involving category-based induction under uncertainty. More important, in some situations, these processes produce systematically different results.

We would argue that our measure of induction is implicit, in the sense the term is often used, for the following reasons. First, people are trained in categories and not in directions; second, the task itself is a same-different task and direction is never mentioned in the test; and third, when people are asked verbally where the objects will go (Chen et al., 2014) or where they think they would look (Experiment 2), they give very different answers from actual motor or eye movements. Recent criticisms of research that suggests that implicit processes improve higher-level cognition relative to explicit processes have noted that many implicit tasks are not actually implicit, as they are consciously accessible when people are queried under proper conditions (see Newell & Shanks, 2014) and "implicit" advantages disappear when appropriate explicit control tasks are used (Newell, 2015; Newell & Shanks, 2014). Thus, it is worth noting that we do not use a criterion of awareness and do not deny that under some conditions subjects may be able to explicitly state how they performed our implicit induction tasks. Rather we focus on different predictions people make with different tasks. Although skeptics may not agree with the term "implicit," it is clear that these motor responses give different inductions from more overt predictions.

Our distinction between implicit and explicit processes in this article may be more reminiscent of Newell and Shanks's (2014) distinction between *deliberative* and *intuitive* processes. We have suggested that our implicit induction tasks, like the motor control tasks that they were inspired by, are more associative and do not elicit the more complex reasoning strategies that our explicit tasks do. Similarly, Newell and Shanks (2014) suggest that intuitive processes, while consciously accessible, are faster and appear more effortless because they rely on learned associations whereas deliberative ones feel more difficult and we seem more aware of them because they are not accompanied by cue-outcome associations (or these associations are more difficult to access). It may be that the form our explicit induction tasks take shifts people into a more deliberative form of reasoning.

Alternative Explanation

An alternative explanation challenges the interpretation that the data from Experiment 1 are a result of integration of information across categories, proposing that the shift toward the alternative category was the result of some sort of practice effect. During the first two phases of the experiment, subjects most often saw the critical shapes move toward the direction of the alternative category. That is, 8 out of the 12 critical shapes went in the same direction that the shapes in the alternative category moved (see Table 1). Perhaps they became used to following the trajectory of the shapes. At test, their eyes may have simply followed the habitual direction of each shape. This seems unlikely for three reasons. First, because the first two phases of Experiments 1 and 2 were identical, subjects in Experiment 2 would have also been trained to follow the direction of each shape. Thus, if the results of Experiment 1 were simply because of a practice effect, subjects in Experiment 2 should have also have shifted their fixations toward the alternative category's direction. They clearly did not. Second, this practice effect does not explain the difference found in the lower and upper visual hemifield (i.e., subjects shifted significantly more toward the direction of the alternative category when that direction was in the upper visual hemifield). That result seems explicable on the basis of subjects shifting fixations more toward the less acute hemifield to better see the stimulus. If they were simply repeating prior tracking behavior, then they would have done so for both hemifields. Third, Experiment 4 showed that shape-direction associations did not in fact cause subjects' eye movements. Those associations were identical to those of Experiment 1, yet there was no evident shift in eye movements.

Role of Explicit Strategies

We began this article by asking why implicit induction may lead to more normative use of category information when categorization is uncertain. We suggested that explicit reasoning is subject to higher-level reasoning strategies that bias people to consider only one category (or possibility) at a time (Evans, 2007), and the implicit inductions avoided such strategies. In these studies, we aimed to ensure that the predictions produced from our implicit task were not based on explicit predictions. First, in our implicit induction task, the prediction was only incidental to the task, making it unlikely that subjects would have consciously been thinking about making a prediction or about how a shape's categorization might impact that prediction. Second, in Experiment 2, we found that subjects seemed to have no explicit knowledge about their actual strategy, suggesting that the integration found in Experiment 1 was not a result of an explicit strategy about where to look. This is consistent with past results showing that people often do not know where they are going to look or have looked (e.g., Theeuwes et al., 1998). Studies of eye movements in skilled squash players shows that they fixate locations where the ball is about to pass through and that they visually pursue the ball's trajectory in a way that suggests that they are predicting where it will go (Hayhoe, McKinney, Chajka, & Pelz, 2012). Those of us who have played racket sports confess to having no such (conscious) strategy of eye movements or even to have ever thought about where to look while playing squash.

The results of Experiment 3 further explain why implicit inductions are sometimes more normative than explicit ones. When subjects were prompted for a categorization before making a speeded implicit induction, they based their predictions on only a single category—the category that they had chosen for their categorization decision. The categorization decisions apparently promoted the disregard of information from alternatives during induction, consistent with the Singularity Principle. These results are particularly interesting as they show that explicit categorizations can influence implicit predictions or at least predictions made in a variety of response modes (as also found in a perceptual judgment task, Stocker & Simoncelli, 2007). Recall our initial example of the person watching an object of unknown categorization falling toward the floor. Our results suggest that simply making a judgment about what category the object most likely belongs to—plates or Frisbees—might greatly influence her decision of what action to take (i.e., staying put or trying to catch it).

A similar explanation may apply to decision-making tasks. As discussed in the Introduction, subjects are sometimes more optimal when executing speeded visuomotor tasks than when making the

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

same decision in an explicit, verbal manner (Trommershäuser et al., 2006, 2008). In the visuomotor tasks, subjects are simply prompted to act quickly, and their prediction about which decision maximizes rewards is only implicit in their action. Indeed, it would be interesting to see if the visuomotor decision making tasks that show near optimal behavior would be subject to similar biases if an explicit judgment were made before the decision making task.

Another salient difference between the implicit and explicit induction tasks that we did not examine is feedback. For both the same-different task from Experiment 1 and the catching task used in Chen et al. (2014), subjects received feedback about whether they performed the task successfully, perhaps thereby teaching them that use of multiple categories would lead to greater success. In the explicit prediction tasks, subjects got no feedback on their predictions. However, there is reason to believe that this is not a critical factor in explaining why the two types of induction lead to different predictions. The block analysis from Experiment 1 (and a similar analysis for the catching task in Chen et al., 2014) found that there was no difference in the predictions in the first and second blocks of test, suggesting that subjects' integrated information across categories early in the test phase. Thus, in these tasks, it seems that feedback-based learning was not a major driver of the integration of information across found in implicit inductions.

Bayesian Models of Higher Order Cognition

Although findings in the category-based induction literature that people use only a single category in induction are not consistent with the claim that much of higher-level cognition can be explained by Bayesian principles, the results of Experiment 3 may be more in line with recent attempts to move Bayesian approaches toward the process level with stochastic sampling algorithms borrowed from machine learning (Goodman et al., 2008; Sanborn, Griffiths, & Navarro, 2010; Vul, Goodman, Griffiths, & Tenenbaum, 2014). The general idea is that the mind may approximate Bayesian inference by evaluating (sampling) a small set of the possible hypotheses in a manner that would reflect the distribution of hypotheses in the posterior (i.e., more likely hypotheses are more likely to be sampled). Depending on the number of samples, such processes can come close to Bayesian inference (and in the limit can approximate optimal inference; Sanborn et al., 2010). From this perspective, ideal Bayesian behavior "emerges only in the average over many learners," as individuals each sample one or a few hypotheses rather than integrating across all samples (Vul et al., 2014, p. 6). The results of Experiment 3, in which people made predictions consistent with their initial categorizations (that were generally probability matched to the true distribution), are consistent with this view. When averaging predictions across all trials, regardless of categorization, these predictions may reflect something like the distribution that would be expected from a Bayesian analysis. However, individual trials do not reflect the use of multiple categories on this account.

Similar patterns have been found in studies of sampling behavior. For example, Denison et al. (2013) asked young children to predict a feature (color) of a toy that came from one of two bags of toys. The distribution of colors was different in each bag. Denison et al. (2013) found that, when aggregated across subjects, predictions reflected an integration of the prior probability that the toy came from each source with distributional information about the colors in each bag. Our probability matching explanation of how subjects made inductions in Experiment 3 (i.e., that they made predictions based on their initial categorization only, but switched between the two possible categorizations over trials), could account for these results. That is, children could be spontaneously picking one bag or the other as the source of the toy and then making their prediction based on only the distribution of colors in that bag. As long as the initial selections of the bag reflects the true base rate of the bags, then this process (aggregated over trials) comes to the same outcome as a full Bayesian analysis. Thus, the children's predictions on any given trial would not truly reflect integration of information, but overall would appear Bayesian.

Our account is similar to the sampling account provided in Denison et al. (2013). However, our claim is that subjects are generally not integrating across hypotheses or categories, while Denison et al. (2013) appeal to Bayesian accounts of behavior and suggest that sampling is part of a "rational strategy for inductive inference" (p. 285), which seems to imply integration over hypotheses. Thus, as theories that rely on sampling algorithms develop, it will be important to clarify whether these accounts of sampling claim that people integrate information across hypotheses, or whether they rely on aggregating trials in which a single hypothesis is considered at a time, probability matched to the distribution of the hypotheses.

Implications for Real-World Inductions

Our results suggest that asking people what category or outcome is most likely may cause them to ignore relevant information and lead to suboptimal predictions even when the predictions are implicit. An analogous problem has been documented in the medical decision making literature. *Diagnosis momentum* is a phenomenon in which, at the expense of possible alternative diagnoses, an initial uncertain diagnosis (i.e., categorization of a patient) is treated as certain in treatment planning (i.e., prediction of how a patient will respond to possible treatments; Croskerry, 2003). Much like the case of medical diagnosis, planning for the future in everyday life often occurs under uncertainty. We may have to pack for a vacation before we know what the weather will be like. Families must save for future educational costs far in advance of knowing what school their children will attend. Our results suggest that, in thinking about future possibilities, it is important to think about multiple possibilities and not to single one out (i.e., ask "What schools do you think little Molly might go to 17 years from now?" rather than the much simpler, "What school do you think little Molly will end up at in 17 years from now?"). Murphy et al. (2012) found that simply asking people to rate the probability of multiple categories increased the likelihood that they integrated information across categories. However, drawing attention to a category can also backfire: Merely rating the likelihood an object was in one of the categories greatly increased people's use of that category in induction, even when the rated likelihood was very low (under 5%, Murphy et al., 2012).

Thus, use of categories is a double-edged sword. It can carry valuable information that allows one to generalize to a new object. However, people may also use categories in a way that limits the amount of information available. Apparently, such problems are more likely when people explicitly evaluate the object's category membership, because doing so can lead to ignoring other possible categories in both explicit and implicit predictions.

Conclusion

We have presented evidence that implicit induction allows for more normative use of category information because these responses avoid reasoning strategies that bias people toward using only a single category. In particular, such responses may allow for integration of information across categories because they avoid explicit categorization of items with uncertain categorization. When people explicitly classify items first, even though the classification is uncertain, our findings suggest that they overcommit to the categorizations, leading to suboptimal use of category information in induction and decision making. The results of the current research extend prior results by demonstrating that explicit categorizations have a strong impact on predictions made in very different response modes. Thus, it will be important for future research to examine the role of identifying only the most likely category or outcome on a wider range of induction and decision making tasks.

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98,* 409–429. http://dx.doi.org/10.1037/0033-295X.98.3.409

Chen, S. Y., Ross, B. H., & Murphy, G. L. (2014). Implicit and explicit processes in category-based induction: Is induction best when we don't think? *Journal of Experimental Psychology: General, 143,* 227–246. http://dx.doi.org/10.1037/a0032064

Croskerry, P. (2003). The importance of cognitive errors in diagnosis and strategies to minimize them. *Academic Medicine, 78,* 775–780. http://dx.doi.org/10.1097/00001888-200308000-00003

Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The Sampling Hypothesis. *Cognition, 126,* 285–300. http://dx.doi.org/10.1016/j.cognition.2012.10.010

Edgar, G. K., & Smith, A. T. (1990). Hemifield differences in perceived spatial frequency. *Perception, 19,* 759–766. http://dx.doi.org/10.1068/p190759

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433. http://dx.doi.org/10.1038/415429a

Evans, J. S. B. (2007). *Hypothetical thinking: Dual processes in reasoning and judgment*. New York, NY: Psychology Press.

Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science, 8,* 223–241. http://dx.doi.org/10.1177/1745691612460685

Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science, 32,* 108–154. http://dx.doi.org/10.1080/03640210701802071

Griffiths, O., Hayes, B. K., & Newell, B. R. (2012). Feature-based versus category-based induction with uncertain categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 576–595. http://dx.doi.org/10.1037/a0026038

Griffiths, O., Hayes, B. K., Newell, B. R., & Papadopoulos, C. (2011). Where to look first for an explanation of induction with uncertain categories. *Psychonomic Bulletin & Review, 18,* 1212–1221. http://dx.doi.org/10.3758/s13423-011-0155-0

Haruno, M., Wolpert, D. M., & Kawato, M. (2001). Mosaic model for sensorimotor learning and control. *Neural Computation, 13,* 2201–2220. http://dx.doi.org/10.1162/089976601750541778

Hayes, B. K., & Chen, T.-H. J. (2008). Clinical expertise and reasoning with uncertain categories. *Psychonomic Bulletin & Review, 15,* 1002–1007. http://dx.doi.org/10.3758/PBR.15.5.1002

Hayes, B. K., & Newell, B. R. (2009). Induction with uncertain categories: When do people consider the category alternatives? *Memory & Cognition, 37,* 730–743. http://dx.doi.org/10.3758/MC.37.6.730

Hayhoe, M. M., McKinney, T., Chajka, K., & Pelz, J. B. (2012). Predictive eye movements in natural vision. *Experimental Brain Research, 217,* 125–136. http://dx.doi.org/10.1007/s00221-011-2979-2

Kahneman, D. (2011). *Thinking fast and slow*. New York, NY: Farrar, Straus & Giroux.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology, 55,* 271–304. http://dx.doi.org/10.1146/annurev.psych.55.090902.142005

Kruglanski, A. W., & Gigerenzer, G. (2011). Intuitive and deliberate judgments are based on common principles. *Psychological Review, 118,* 97–109. http://dx.doi.org/10.1037/a0020762

Lagnado, D. A., & Shanks, D. R. (2003). The influence of hierarchy on probability judgment. *Cognition, 89,* 157–178. http://dx.doi.org/10.1016/S0010-0277(03)00099-4

Levine, M. W., & McAnany, J. J. (2005). The relative capabilities of the upper and lower visual hemifields. *Vision Research, 45,* 2820–2830. http://dx.doi.org/10.1016/j.visres.2005.04.001

Malt, B. C., Ross, B. H., & Murphy, G. L. (1995). Predicting features for members of natural categories when categorization is uncertain. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 646–661. http://dx.doi.org/10.1037/0278-7393.21.3.646

Murphy, G. L., Chen, S. Y., & Ross, B. H. (2012). Reasoning with uncertain categories. *Thinking & Reasoning, 18,* 81–117. http://dx.doi.org/10.1080/13546783.2011.650506

Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology, 27,* 148–193. http://dx.doi.org/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. http://dx.doi.org/10.1016/j.cognition.2004.01.009

Murphy, G. L., & Ross, B. H. (2010a). Category vs. object knowledge in category-based induction. *Journal of Memory and Language, 63,* 1–17. http://dx.doi.org/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. http://dx.doi.org/10.1037/a0018685

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature, 434,* 387–391. http://dx.doi.org/10.1038/nature03390

Najemnik, J., & Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision, 8,* 1–14. http://dx.doi.org/10.1167/8.3.4

Newell, B. R. (2015). "Wait! Just let me not think about that for a minute." What role do implicit processes play in higher-level cognition? *Current Directions in Psychological Science, 24,* 65–70. http://dx.doi.org/10.1177/0963721414551958

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. http://dx.doi.org/10.3758/PBR.17.6.869

Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences, 37,* 1–19. http://dx.doi.org/10.1017/S0140525X12003214

Papadopoulos, C., Hayes, B. K., & Newell, B. R. (2011). Noncategorical approaches to feature prediction with uncertain categories. *Memory & Cognition, 39,* 304–318. http://dx.doi.org/10.3758/s13421-010-0009-4

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

Rosch, E. H. (1973). Natural categories. *Cognitive Psychology, 4,* 328–350.

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. http://dx.doi.org/10.1037/0278-7393.22.3.736

Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review, 117,* 1144–1167. http://dx.doi.org/10.1037/a0020511

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3–22. http://dx.doi.org/10.1037/0033-2909.119.1.3

Stanovich, K. (2009). *What intelligence tests miss: The psychology of rational thought*. New Haven, CT: Yale University Press.

Stocker, A. A., & Simoncelli, E. P. (2007). A Bayesian model of conditioned perception. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in Neural Information Processing Systems 20: 21st Annual Conference on Neural Information Processing Systems 2007* (pp. 1409–1416). Norwich, England: Curran Associates.

Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *The Journal of Neuroscience, 26,* 10154–10163. http://dx.doi.org/10.1523/JNEUROSCI.2779-06.2006

Theeuwes, J., Kramer, A. F., Hahn, S., & Irwin, D. E. (1998). Our eyes do not always go where we want them to go: Capture of the eyes by new objects. *Psychological Science, 9,* 379–385. http://dx.doi.org/10.1111/1467-9280.00071

Trommershäuser, J., Körding, K. P., & Landy, M. S. (Eds.). (2011). *Sensory cue integration*. New York, NY: Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780195387247.001.0001

Trommershäuser, J., Landy, M. S., & Maloney, L. T. (2006). Humans rapidly estimate expected gain in movement planning. *Psychological Science, 17,* 981–988. http://dx.doi.org/10.1111/j.1467-9280.2006.01816.x

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, 20,* 1419–1433. http://dx.doi.org/10.1364/JOSAA.20.001419

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences, 12,* 291–297. http://dx.doi.org/10.1016/j.tics.2008.04.010

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. http://dx.doi.org/10.3758/BF03193065

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science, 38,* 599–637. http://dx.doi.org/10.1111/cogs.12101

Wickens, C. D., & Hollands, J. G. (2000). *Engineering Psychology and Human Performance*. Upper Saddle River, NJ: Prentice Hall.

Zhu, J., & Murphy, G. L. (2013). Influence of emotionally charged information on category-based induction. *PLoS ONE, 8,* e54286. http://dx.doi.org/10.1371/journal.pone.0054286

Received April 23, 2014
Revision received October 13, 2015
Accepted October 18, 2015 ■

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.