

Shape memory for intrinsic versus accidental holes

ROLF NELSON

Wheaton College, Norton, Massachusetts

AND

JESSICA THIERMAN AND STEPHEN E. PALMER

University of California, Berkeley, California

Two experiments were performed to investigate the circumstances under which the shape of a visual hole is perceived and remembered. A distinction is made between an *intrinsic hole*, a surrounded region that is bordered on all sides by a single surface at a continuous depth, and an *accidental hole*, a surrounded region with borders belonging to multiple surfaces at discontinuous depths. In the first experiment, we demonstrated a marked memory advantage for intrinsic holes over both accidental holes and accidental parts when the accidental hole and part were nested within the intrinsic hole. Indeed, the accidental holes and parts appeared not to be remembered at all. In Experiment 2, we used nonnested display sets that ruled out the possibility that the results from the first experiments were due to any image-based differences and included solid objects, as well as intrinsic and accidental holes. The results showed that the shapes of intrinsic holes were still remembered better than those of accidental holes, but not as well as those of solid objects. The results are discussed in a framework of ecological validity, arguing that observers are likely to encode shapes that are invariant features of the environmental scene and not those that arise only from the coincidence of particular viewing conditions.

Recently, there has been much interest in visual holes as a tool in understanding figure-ground and depth relations across an edge (Bertamini, 2006; Bertamini & Croucher, 2003; Bertamini & Farrant, 2006; Bertamini & Mosca, 2004; Casati & Varzi, 1994; Hulleman & Humphreys, 2005; Nelson & Palmer, 2001; Palmer, 1999; Palmer, Davis, Nelson, & Rock, 2008). One of the more intriguing questions that has been asked about visual holes is whether or not their interior shapes are perceived. From a classical gestalt viewpoint (e.g., Rock, 1983; Rubin, 1921), the process of figure-ground organization imparts border ownership, in an all-or-nothing fashion, to one side of a contour or the other. The side that is perceived as the “figure” is seen as both closer in depth and shaped by the contour, whereas the side that is perceived as “ground” is seen as continuing behind the figure, and not being shaped by it. However, a surrounded region that is perceived to be a hole appears to violate this coupling, since the shape of the interior region is perceived and remembered, even though it is perceived as being a farther ground, in terms of depth (Palmer et al., 2008).

Following Palmer et al. (2008), we define a *visual hole* as a closed image contour, every point of which is the projection of an environmental point that is closer to the observer than the projection of its adjacent interior point. This definition is merely the starting point in the perceptual organization of an image that might eventually result

in the perception of a physical hole in the environment. The term *visual hole* is thus not intended to be synonymous with a *physical hole* in the environment. The problem faced by the visual system is therefore how to get from an image-based visual hole to the perception of a corresponding physical hole in the environment.

All of the visual holes investigated in recent studies are what we will refer to as *intrinsic holes*: holes with boundaries that are environmentally continuous and constitute the inner closed contour of a single surrounding surface or object. We call such holes *intrinsic* because they are invariant properties of the material extension of the hole-bearing objects. The shapes of intrinsic holes have been found to be remembered as well as those of solid objects in previous research using relatively simple displays (Palmer et al., 2008).

In contrast, there are other sorts of visual holes that are not intrinsic features of any single surface or object. We call these *accidental holes* because they are determined by coincidences of perspective projection in which the projected images of multiple objects or surfaces overlap to form a visual hole from that particular viewpoint. Accidental holes are examples of visual holes, because their entire bounding image contour consists of environmental points that lie closer to the observer than do their adjacent interior points, as defined above. However, they differ from intrinsic visual holes in that the bounding contour

R. Nelson, rnelson@wheatonma.edu

in the image does not correspond to a closed contour in the 3-D environment, but to one in which there are depth discontinuities at T-junctions along its perimeter.¹ Because accidental holes do not constitute invariant properties of any single surface or object, but rather the result of a coincidental alignment of two or more objects from some particular point of view, we expect their shape to be poorly perceived and poorly remembered. Findings of this sort have been reported previously in the literature. For example, He and Nakayama (1992) found longer visual search times for an “L” when part of the contour forming it was owned by an occluding surface, but they did not study displays involving accidental holes.

Figure 1 shows an image in which both an intrinsic hole and an accidental hole are present. The dark gray object contains an intrinsic hole that is an integral part of it, in the sense that its entire perimeter is a closed, connected subset of the boundary of the single surrounding object. The other hole in the image is an accidental hole formed partly by the boundary of the intrinsic hole in the closer, dark gray object and partly by the boundary of the farther, light gray object, a portion of which is visible through the intrinsic hole. The accidental hole is not intrinsic, because its perimeter is not a closed contour in the environment: There are clear depth discontinuities at the T-junctions along its perimeter. Moreover, there are viewpoints from which this accidental hole simply would not exist as a visual hole in the projected image of the same scene (i.e., from any oblique view in which the contours of the two objects either do not overlap or overlap in ways that do not reveal part of the farther surface through the nearer object’s intrinsic hole).

The distinction between intrinsic and accidental holes is theoretically important because it constrains possible answers to questions about why the shapes of intrinsic holes might be well perceived and well remembered, as Palmer et al. (2008) found them to be, despite the ground status of their interiors. One possibility is that the visual system encodes the shape of every visual hole, perhaps as a result of encoding the shape of every 2-D region in the image.

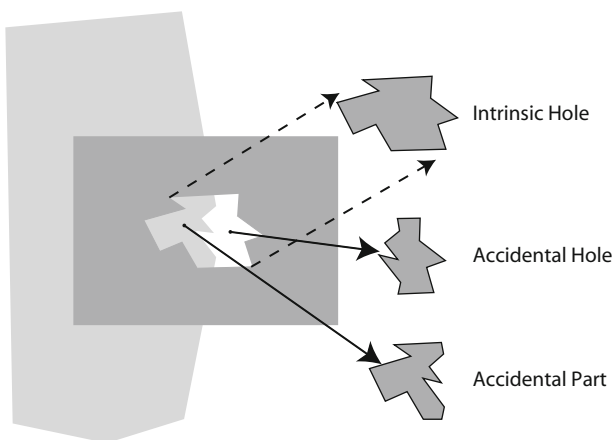


Figure 1. An example of the displays used in Experiment 1. Each image contained a rectangular object in front, with a hole in it, and an object behind, which was partially visible through this hole. The three shapes identified in the figure were later tested for memory.

Without further assumptions, however, this would predict no difference between intrinsic and accidental holes in terms of shape encoding. Another possibility is that the shapes of visual holes are remembered only if they are encoded as integral and invariant aspects of the shapes of the objects that contain them. This view implies that the shapes of intrinsic holes should be remembered, whereas those of accidental holes should not.

EXPERIMENT 1

Memory for Intrinsic Versus Accidental Holes in Nested Displays

In the first experiment, we presented to participants stereoscopic depth images like the one depicted in Figure 1 to identify any differences in the perception of and memory for the shapes of intrinsic versus accidental holes. Each image consisted of two flat, overlapping, 2-D objects in depth (one green, one red), the farther of which was solid and the closer of which contained an intrinsic hole. Through the intrinsic hole of the closer surface, the observer could see a portion of the background surface behind both objects (the accidental hole) and a portion of the farther object (which we term, by analogy, the *accidental part*). We call these displays *nested*, because the intrinsic hole is the union of the accidental part and accidental hole. After seeing a series of such displays, participants were tested for their recognition memory of the shapes of the regions previously presented as intrinsic holes, accidental holes, and accidental parts, together with similar novel distractors. By the logic described above, we expected that memory for the intrinsic holes would be greatly superior to memory for either the accidental holes or the accidental parts, which we expected to be at or near chance. That is, we believed that the shapes of accidental holes would not be encoded or remembered because they are not invariant properties of objects, whereas the shapes of intrinsic holes would be encoded and remembered because they are.

Method

Participants. Twenty-four undergraduates at the University of California, Berkeley, volunteered to participate in the experiment for partial course credit in psychology courses.

Displays. A set of 24 displays like the ones in Figure 1 were constructed, each of which consisted of two overlapping shapes. The closer shape was a rectangle with a hole cut through it that revealed a portion of the farther shape underneath and the farthest background surface. The intrinsic hole in the rectangle was the union of two sub-regions of equal area—an accidental part and an accidental hole—both of which were bounded by a portion of the border of the intrinsic hole and by a portion of the border between the farther object and the background surface. The displays were left–right counterbalanced, so that half of the displays were mirror reversals. In the presentation phase, each participant saw 3 displays with the background object on the left and 3 with the background object on the right.

The displays were seen through a stereoscopic viewer consisting of a hood placed over the computer monitor that separated the screen into two images. Stereoscopic depth was induced by laterally displacing the background shape relative to the foreground shape, which necessarily introduced a small difference in the position of the border between the images in the left-eye and right-eye displays. In order to be sure that the perceived shapes of the regions seen in the presentation phase (which were viewed stereoscopically) were simi-

lar to those in the memory test phase (which were not viewed stereoscopically), two precautions were taken. First, the borders of the intrinsic hole near the center, where the border was positioned, were always straight and approximately horizontal in order to eliminate qualitative geometrical differences between the left and right views of the accidental part and accidental hole. Second, each participant was tested to determine his or her dominant eye, with respect to stereoscopic viewing. Two entire sets of test displays were then constructed—one for participants who were right-eye dominant, and one for those who were left-eye dominant. These displays matched the view the participants saw with their dominant eye. This ensured that the shapes experienced during the test sequence were maximally similar to those experienced during the initial presentation.

The regions used in the memory test were the intrinsic hole, the accidental hole, and the accidental part for each display. They were presented individually as gray regions on a homogenous blue background. Because the sizes of the three types of test shapes were different in their initial presentation, all test displays were presented at an intermediate size that was approximately midway in area between the size of the intrinsic hole and of the accidental parts and accidental holes. Thus, the sizes of the test figures were never the same as in the presentation phase, being bigger for the accidental test items and smaller for the intrinsic test items.

Procedure. Participants were first tested for eye dominance by having them make an aperture with their hands at arm's length and use it to sight an object with both eyes open. They were then asked to close each eye separately and to determine which eye's view contained the object they initially saw through their hands when both eyes were open. Their dominant eye was assessed to be the one for which its monocular image was most similar to the binocular image.

Participants were then shown the stereoscopic viewer and given the following instructions: "You will be shown a series of six pictures that will be repeated three times each. Your only task is to study these pictures, as there will be a memory test later in the experiment." Each of the displays remained on the screen for 4 sec. The order of presentation of the displays was counterbalanced factorially across participants.

Afterward, the stereoscopic viewer was taken off the screen, and the participants were told that they would now be tested for their memory of the shapes they had seen. The gray shapes were displayed for 1.5 sec each, after which participants were instructed to press "O" if it was an old region (one they had seen previously), and "N" if it was a new region (one they had not seen before). After this, they were instructed to give a rating of 1 (*not sure*), 2 (*somewhat confident*), or 3 (*very confident*) to indicate how confident they were in each response.

In the test phase, exactly half of the regions were old; the other half were new. Participants saw all 18 regions from the presentation phase (the 6 examples each of intrinsic holes, accidental holes, and accidental parts), plus 18 regions from a set of similar displays not seen by that participant in the presentation phase.

Results

The results showed that the shapes of the intrinsic holes were remembered much better than those of the accidental holes and parts. A standard signal detection measure (d') was computed for each region type, aggregated across subjects, using correct responses to old displays as hits and incorrect responses to new displays (as being old) as false alarms. These d' scores were much greater for intrinsic holes ($d' = 1.18$) than for either accidental holes ($d' = 0.14$) or accidental parts ($d' = 0.25$).²

Confidence ratings corroborated these results. These ratings were a reliable indicator of memory; the rating was highly correlated with the percentage of correct answers ($r^2 = .97$). The data were coded so that each confidence rating was signed positively if it followed an *old* response and negatively if it followed a *new* response. Figure 2




Figure 2. Results of Experiment 1. Intrinsic holes were remembered with a much higher confidence than were accidental holes or parts.

shows the average confidence ratings. Of the three types of test regions, only the old intrinsic holes were confidently remembered as having been seen before, with a mean rating of +1.54, which was significantly higher than that for accidental holes (-0.73) [$t(23) = 8.87, p < .001$] or that for accidental parts (-0.48) [$t(23) = 8.04, p < .001$]. In comparison with their having seen the new shapes, participants were much more confident that they had seen the intrinsic holes [$t(23) = 8.32, p < .001$].

Previously seen accidental holes were not rated differently than their corresponding new shapes were [with a mean rating of -0.94 for new shapes vs. -0.73 for previously seen shapes, $t(23) = 1.00$]. However, previously seen accidental parts were rated as marginally higher than their corresponding new shapes [a mean rating of -0.97 for new vs. -0.48 for previously seen shapes, $t(23) = 2.18, p < .05$]. Although this is a small difference, the slightly better performance for accidental parts than for accidental holes may reflect the fact that the accidental parts contained an intrinsic contour along the side that belonged to the amodally completed object partly seen through the intrinsic hole, whereas the accidental holes contained no intrinsic contours at all.

We attribute the striking differences in memory for the shapes of intrinsic versus accidental holes to the fact that the shapes of intrinsic holes are invariant properties of single objects, whereas accidental holes are highly variable properties of particular projections due to the particular arrangement of two or more objects relative to the viewpoint. Other explanations of the results must be considered, however.

In this set of nested displays, where the accidental hole was always inside the intrinsic hole, there were unavoidable differences between these two regions in terms of both size and shape. Because the intrinsic holes were the geometric union of the accidental holes and accidental parts (see Figure 1), the intrinsic holes were necessarily about twice as large as either of the other two regions, which were about the same in size. Even though shape perception is generally considered to be invariant over size transformations (see Palmer, 1999, p. 364), an objection could be raised that this difference in size was the reason for the large differences in memory performance. People may simply remember

larger regions better than smaller ones, for example. Alternatively, people may be able to remember a shape better when the size of the test region is smaller than the original than they do when it is larger than the original. Perhaps these or other size effects are reflected in the data.

Other potential confounds concern differences in the shapes of the intrinsic and accidental holes and their nested structure. Both result from the fact that the intrinsic holes were the union of the accidental hole and accidental part regions. Because the latter two regions were roughly isotropic, the intrinsic holes were necessarily more elongated than the accidental regions. Perhaps people simply remember elongated shapes better than isotropic ones, and this difference caused the difference in shape memory between intrinsic holes and the accidental regions. Another possibility is that people simply remember the union of two adjacent regions better than they remember the component regions themselves. Experiment 2 addressed all of these possible confounds by eliminating them in the display materials.

EXPERIMENT 2

Memory for Objects and Holes in Nonnested Displays

Experiment 2 had two main purposes. First, it eliminated the physical differences between the intrinsic and accidental hole regions in Experiment 1 that were just discussed (i.e., differences in size, shape, and nesting). Second, it included solid objects in the perceptual conditions, so that a specific comparison could be made between people's memory for objects, intrinsic holes, and accidental holes. Using much simpler displays, Palmer et al. (2008) found no difference between memory for the shapes of holes and those of solid objects, and we wanted to find out whether this result would replicate in the present, more complex displays. A further improvement introduced in Experiment 2 was the use of LCD shutter glasses to create a more realistic perception of the shapes in 3-D space.

Method

Participants. Sixteen students from Wheaton College volunteered for the experiment. They received compensation in the form of a \$5 gift certificate to the campus bookstore.

Apparatus. During the presentation phase, images were presented on a 19-in. computer monitor with a pair of Stereographics CrystalEyes LCD shutter glasses. The monitor ran at a resolution of 800×600 pixels with a refresh rate of 140 Hz, providing a frame rate to each eye of 70 Hz. The three depth planes were depicted at approximately 0° , 0.8° , and 1.2° of lateral disparity from the viewpoint of the observers, who were approximately 20 in. from the monitor and not constrained by a chinrest.

Each image contained two of the three regions of interest: an object, an intrinsic hole, and/or an accidental hole. These regions were embedded in a stereoscopic display depicting the three depth planes defined above (see Figure 3). The farthest plane was always a randomly textured background to ensure good stereoscopic fusion. The closer two planes contained either two or three flat surfaces, some of which were solid and some of which contained an intrinsic hole. Other than the textured background, each display contained four regions, each of which corresponded to (1) a solid flat surface that was fully visible (an object); (2) a flat surface that contained an intrinsic hole, through which the textured ground plane was visible (an intrinsic hole); or (3) a flat surface that overlapped with another flat surface to create an

accidental hole, through which the textured ground plane was visible (an accidental hole). Figure 3 shows an example of one such quartet of images, using cast shadows rather than binocular disparity to portray depth relations, each image (without the shadows) being half of a stereoscopic pair. Figure 3A shows an intrinsic hole on the left and an object on the right. Figure 3B shows an accidental hole on the left and an intrinsic hole on the right. Figure 3C shows an accidental hole on the left and an object on the right. Figure 3D shows an accidental hole on the left and an intrinsic hole on the right (the same as in Figure 3B, but with the depth planes reversed). The geometry of the quartets was such that the shapes of the intrinsic holes were sometimes the same as those of the solid objects (Figures 3B and 3D vs. 3A and 3C) and sometimes the same as those of the accidental holes (Figure 3A vs. 3B, 3C, and 3D), but the shapes of the solid objects and accidental holes were never the same.

There were eight different such quartets, for a total of 64 images. The background was black-and-white noise, and the solid regions were orange and blue. These combinations were presented in random order to participants in a Latin squares design, so that each participant saw only one display from each quartet of images, as defined above, for a total of eight displays per participant.

The test shapes were all solid gray in color and were presented against a white background. There were 32 test shapes, 16 of which were old (2 in each of the eight 3-D images that they had viewed in the initial presentation phase) and 16 of which were new shapes, different from, but similar to, ones the participants had already seen.

Procedure. Before the experiment, all participants were tested for stereo vision by viewing a random-dot stereogram via a stereo viewer and being asked what shape they saw in depth. All participants correctly reported seeing a square. Participants were also tested for color vision using a set of Ishihara plates, and all were found to have normal color vision. As in Experiment 1, they were tested for their dominant eye, and the set of test displays was selected accordingly.

The participants then put on the LCD shutter glasses and were instructed that they would see a sequence of eight 3-D images for 4 sec each, and that this set of images would be repeated three times. They were told to remember the images, because their memory would later be tested.

After the participants had seen the presentation images, they were told that they would see a series of gray shapes, and that they were to indicate whether or not they had seen the shape before in the 3-D images. They were to press a button on the keyboard (which was overlaid with labels) to indicate how certain they were that they had seen the shape before. The buttons were labeled, left to right, -3 , -2 , -1 , $+1$, $+2$, and $+3$, and the participants were told that -3 meant they were very sure they had not seen the shape, -2 meant they were somewhat sure they had not seen the shape, -1 meant that they were not sure but thought they had not seen the shape, and so on. They were told to make the best guesses they could.

Participants viewed the test shapes via SuperLab 3.0 data collection software. They were given a total of 32 shapes to judge, 16 of which were new to them and 16 of which they had seen in the presentation images. The order of the test shapes was random, and they appeared on the screen until a response was made.

Results

Consistent with the results of Experiment 1, the present results showed that the shapes of intrinsic holes were remembered better than those of accidental holes. An aggregate d' across all participants was measured by collapsing the confidence responses into *old* ($+1$ to $+3$) and *new* responses (-1 to -3). Discrimination was highest for objects ($d' = 2.12$), next highest for intrinsic holes ($d' = 1.26$), and poorest for accidental holes ($d' = 0.65$).

Figure 4 shows average confidence ratings by region type. Again, these ratings correlated highly with correct responses, so that higher confidence ratings were associated

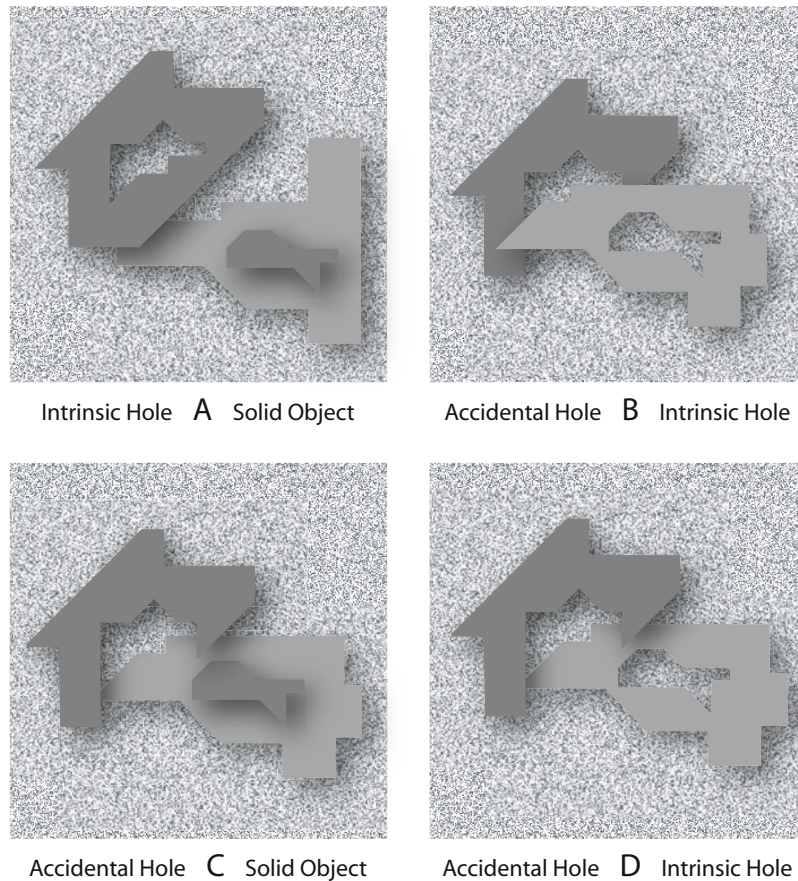


Figure 3. Displays used in Experiment 2. (See text for details.)

with higher probabilities of correct responses ($r^2 = .99$). The average rating for accidental holes was -0.31 , indicating that participants tended to perceive them as (relatively) new, whereas the average rating for intrinsic holes was $+0.84$, indicating that participants tended to perceive them as (relatively) old. This difference was highly significant [$t(15) = 4.15, p < .001$]. Confidence ratings for intrinsic holes were reliably lower than those for objects, however [$t(15) = 4.60, p < .001$]. The present results also showed some memory

for accidental holes, in that the average confidence rating for accidental holes (-0.31) was reliably higher than that for new regions (-1.46) [$t(15) = 4.56, p < .001$].

In order to ensure that these differences in memory were not due to some shapes' being more memorable than others, a specific comparison was made between identical shapes when they were presented as different region types. Consistent with the above results, shapes presented as intrinsic holes were rated as more confidently remembered (0.91) than the identical shapes presented as accidental holes (-0.31) [$t(15) = 3.58, p < .01$]. Shapes presented as solid objects were also rated as more confidently remembered (2.0) than the identical shapes presented as intrinsic holes (0.81) [$t(15) = 2.92, p < .01$].³ These memory differences are thus uncontaminated by any differences in size, shape, or nesting of the regions involved. There are further differences that may be relevant to the pattern of results, however. One such difference is the fact that the solid objects were presented and tested as solid regions, whereas the intrinsic and accidental holes were initially presented as textured regions but tested as solid regions. This may be responsible for objects' being better remembered than holes, but it cannot be responsible for intrinsic holes' being better remembered than accidental holes, because both sorts of holes were initially presented as textured regions.

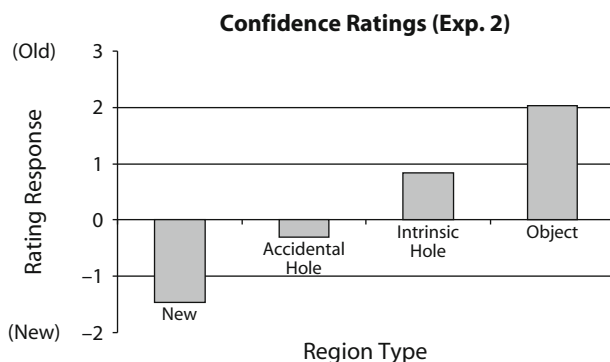


Figure 4. Results of Experiment 2. New shapes and accidental holes were rated as new, whereas intrinsic holes and objects were rated as having been seen before.

DISCUSSION

The results from both of the present experiments support our initial conjecture that the shapes of intrinsic holes, which are surrounded on all sides by a single connected surface, would be remembered better than those of accidental holes, for which the projected shape is bordered by multiple surfaces at different depths. This difference is not specific to the nested displays used in Experiment 1, because a similar difference was found in the nonnested displays of Experiment 2, which also controlled for other possible confounds present in the nested displays of Experiment 1.

Some of the results of Experiment 2 are inconsistent with those reported by Palmer et al. (2008), who found no reliable difference between memory for solid objects and that for intrinsic holes. In contrast, memory performance in Experiment 2 was significantly better for solid objects than for intrinsic holes. Several explanations are possible: One that has already been mentioned is that the objects were both presented and tested as solid regions, whereas the holes were presented as textured regions and tested as solid ones. Another is that although Palmer et al. never found statistically worse memory for intrinsic holes than for objects, both of their experiments showed a trend in this direction. Perhaps the greater memory load with multiple objects and/or holes in the present displays simply amplified this small difference to the point where it was statistically reliable. A third is that attentional effects may be responsible. The displays that Palmer et al. used were much simpler, containing a background surface that was partly occluded by a single connected region perceived either as a solid object or as a square object bearing an asymmetrically curved hole, through which the ground was visible. Thus, on any given trial, only a single object or a single intrinsic hole was present in the display—conditions that plausibly would lead observers to pay equal attention to the object and the hole. The displays in the present Experiment 2 were much more complex, with objects and intrinsic holes probably competing with each other for attention. It is quite possible that the worse memory we found here for intrinsic holes relative to the objects arose from objects' receiving a higher attentional priority than intrinsic holes. It is well established that attended items are better remembered than unattended items (e.g., Rock & Gutman, 1981), so if attention is biased toward encoding objects, the difference we found in Experiment 2 between memory for objects and that for intrinsic holes may be explained by such an attentional strategy.

There is also an inconsistency between the results of Experiment 1 and those of Experiment 2 concerning memory for accidental holes: In Experiment 1, the shapes of the accidental holes were essentially not remembered at all ($d' = 0.14$), whereas in Experiment 2, there was a small but reliable amount of memory ($d' = 0.65$). A possible explanation for this memory difference arises from the structure of the nested versus nonnested displays. In Experiment 1, where the accidental holes were nested within the intrinsic holes, the intrinsic holes had a very different appearance from the accidental ones, the former being larger, more elongated, and divided into two distinct subregions. These differences may have made it very easy

for observers to select the intrinsic holes for attention and further processing without also processing the accidental holes or accidental objects. In Experiment 2, the intrinsic and accidental holes were much more similar in appearance, being the same in size and elongation, and filled with the same noise texture (see, e.g., Figures 3B and 3D). This similarity may have required observers to pay more attention to the accidental holes in Experiment 2 while they were trying to identify which holes were intrinsic (and, therefore, worth further processing) and which were not. This additional attentional processing may have produced better memory for them in the later test trials.

In the foregoing discussion, we have appealed to attentional effects as plausible explanations of differences in performance on these shape memory tasks. Indeed, it is possible that all of the differences we found may be explained by attentional effects. We are suggesting not that observers are necessarily employing conscious attentional strategies, but that prior experience with objects, intrinsic holes, and accidental holes has effectively educated the visual system to deploy attention to those properties that are stable over changes in viewpoint. That is, the visual system may have learned to attend to properties of objects and intrinsic holes, which are viewpoint invariant, and not to accidental holes, which vary greatly with viewpoint. It seems almost beyond question, for example, that in the absence of explicit instructions to do otherwise, most people would encode the displays using a default attentional strategy with a strong priority toward encoding the shapes of objects.

This line of reasoning suggests that the results of Experiment 2 could have been significantly different if participants had received specific instructions about how to attend to the displays. One group of observers could have been asked to attend explicitly to accidental holes in such displays, another to attend explicitly to intrinsic holes, and a third to attend explicitly to objects. Such attentional instructions could even have produced effects that were the reverse of those we found, although there could still have been remaining residual differences between intrinsic and accidental holes. We are in the process of investigating this possibility.

An issue that we will raise but not settle is whether problems of interpretation could have arisen from our using an explicit memory task to draw inferences about the perception of shapes at their initial presentation. There is a long and productive history of using memory performance to assess the outcome of perceptual processing (e.g., Driver & Baylis, 1996; Rock, 1973; Rock & DiVita, 1987; Rock & Engelstein, 1959), and we are not aware of any serious problems that have arisen from this practice. Nevertheless, one could argue that people normally do not perceive the shapes of intrinsic holes, and that our participants did so only because they knew that their memory would later be tested. In any case, an important prediction of this type of objection is that the present results would fail to replicate under incidental learning conditions, such as asking participants to view the displays in order to make aesthetic ratings of them, say, as sculptures, and then to test their memory without prior warning. We doubt that this would greatly change the pattern of results, but we are in the process of finding out.

Before closing, we want to consider the important theoretical question that arises from the results of our experiments: Why are intrinsic holes remembered so much better than accidental holes? In prior research (Palmer et al., 2008), we argued that the shapes of intrinsic holes are remembered as a sort of “immaterial surface”; that is, they are encoded much like their material counterparts, but with an additional “missing” or “empty” tag specifying that they consist of the absence of surface. Furthermore, because this immaterial surface “belongs to” the supporting material surface, the intrinsic hole becomes an aspect of this hole-bearing object. From this perspective, it follows that accidental holes should not be remembered, precisely because there is no object to which they belong. There is simply no shape of an immaterial surface that corresponds to the missing part of an object, because the contours are not integrated as part of a continuous surface—at least not in the representation of the 3-D scene that gave rise to the accidental hole.⁴ This makes sense from an ecological perspective, because accidental holes are not invariant features of any object, but are complex and largely coincidental projective features of the relation between a scene and the observer’s particular viewpoint.

Another way of looking at the present results is in terms of their implications for the classical gestalt view of figure–ground perception. We believe that they actually clarify and support this view rather than contradicting it. To see why, one must first realize that perceiving and remembering the shape of an intrinsic hole’s interior region would violate the coupling of depth and shape perception only if the ground were perceived as being shaped by the shared edge. It is not. Rather, it is the figure that is shaped by the edge of an intrinsic hole—namely, the limit of its extension toward the inside—because the ground is clearly perceived as continuing seamlessly behind the hole’s edge. It therefore makes good sense that the shape of an intrinsic hole’s inner contour is perceived and remembered, and the fact that it reflects the observer’s knowledge of the shape of the figural region, not that of the ground. This is not the case for accidental holes, however. The shape of an accidental hole’s contour is not the shape of any object or surface in the scene, but merely a coincidence of viewpoint. Indeed, that contour does not even exist in the physical world as a continuous entity, because it jumps back and forth discontinuously in depth from one object to another. The surprising thing is that there might be any memory for it at all.

AUTHOR NOTE

This research was supported in part by Grant 1-R01-MH46141 from the National Institute of Mental Health to the third author and Irvin Rock, and by the Wheaton Research Partnership at Wheaton College. We thank Mary Peterson, the members of the PPRI research group at Berkeley, and members of the faculty at Wheaton College for their input and support. Correspondence concerning this article should be addressed to R. Nelson, Psychology Department, Wheaton College, Norton, MA 02766 (e-mail: rnelson@wheatonma.edu).

Note—Accepted by the previous editorial team, when Thomas H. Carr was Editor.

REFERENCES

- BERTAMINI, M. (2006). Who owns the contour of a visual hole? *Perception*, *35*, 883-894.
- BERTAMINI, M., & CROUCHER, C. J. (2003). The shape of holes. *Cognition*, *87*, 33-54.
- BERTAMINI, M., & FARRANT, T. (2006). The perceived structural shape of thin (wire-like) objects is different from that of silhouettes. *Perception*, *35*, 1679-1692.
- BERTAMINI, M., & MOSCA, F. (2004). Early computation of contour curvature and part structure: Evidence from holes. *Perception*, *33*, 35-48.
- CASATI, R., & VARZI, A. C. (1994). *Holes and other superficialities*. Cambridge, MA: MIT Press.
- DRIVER, J., & BAYLIS, G. C. (1996). Edge-assignment and figure–ground segmentation in short-term visual matching. *Cognitive Psychology*, *31*, 248-306.
- HE, Z. J., & NAKAYAMA, K. (1992). Surfaces versus features in visual search. *Nature*, *359*, 231-233.
- HULLEMAN, J., & HUMPHREYS, G. W. (2005). Differences between searching among objects and searching among holes. *Perception & Psychophysics*, *67*, 469-482.
- MALIK, J. (1987). Interpreting line drawings of curved objects. *International Journal of Computer Vision*, *1*, 73-103.
- NELSON, R., & PALMER, S. E. (2001). Of holes and wholes: The perception of surrounded regions. *Perception*, *30*, 1213-1226.
- PALMER, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
- PALMER, S. E., DAVIS, J., NELSON, R., & ROCK, I. (2008). Figure–ground effects on shape memory for objects versus holes. *Perception*, *37*, 1569-1586.
- ROCK, I. (1973). *Orientation and form*. New York: Academic Press.
- ROCK, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
- ROCK, I., & DiVITA, J. (1987). A case of viewer-centered object perception. *Cognitive Psychology*, *19*, 280-293.
- ROCK, I., & ENGELSTEIN, P. (1959). A study of memory for visual form. *American Journal of Psychology*, *72*, 221-229.
- ROCK, I., & GUTMAN, D. (1981). The effect of inattention on form perception. *Journal of Experimental Psychology: Human Perception & Performance*, *7*, 275-285.
- RUBIN, E. (1921). *Visuell wahrgenommene Figuren*. Copenhagen: Glydenal.

NOTES

1. We do not mean that T-junctions necessarily imply the existence of accidental, rather than intrinsic, holes in the general case of 3-D physical holes. A torus (or doughnut), for example, often projects T-junctions at two points along its inner contour from an oblique viewing perspective, even though the visual hole it contains is intrinsic. The reason is that 3-D curved surfaces can be self-occluding (cf. Malik, 1987). In the more restricted case of homogeneously colored, flat surfaces, such as the ones we use here, however, accidental holes do imply the presence of T-junctions along the hole’s contour.

2. The d' results were determined by first averaging the hit rates and false alarm rates across participants, because some participants exhibited 100% hit rates and/or 0% false alarm rates, which implied d' scores of infinity. Because each condition then gives rise to a single d' measure, no statistical analyses could be performed on the data in this form. Standard statistical analyses were carried out on the confidence rating data, however, for which each participant’s average rating could be computed separately in each condition.

3. Recall that these averages here are only for the specific shapes that formed the individual comparisons (e.g., the shapes that were seen as both intrinsic holes and accidental holes), and are thus different from the overall averages for each type of region as described above.

4. We add this caveat about 3-D interpretations because an accidental hole’s shape can always be encoded at a 2-D level, where it actually does exist as a closed contour that has a defined shape.

(Manuscript received March 22, 2007;
revision accepted for publication August 7, 2008.)