Perception, 2002, volume 31, pages 521 – 530

DOI:10.1068/p3223

Experiments on the role of painted cues in Hughes's reverspectives

Thomas V Papathomas
Laboratory of Vision Research and Department of Biomedical Engineering, Rutgers University, Psychology Building, Busch Campus, Piscataway, NJ 08854, USA;
e-mail: papathom@zeus.rutgers.edu
Received 14 March 2001, in revised form 17 December 2001

Abstract. The English artist Patrick Hughes has created an extraordinary class of painted artpieces, most commonly referred to as 'reverspectives'. They consist of truncated pyramids and prisms with their smaller faces closer to the viewer, in such a way as to allow a realistic scene to be painted on them. The works of art contain rich perspective and other painted cues that conspire to elicit an illusory depth percept that is the reverse of the physical depth arrangement. This reverse depth is obtained under a wide range of viewing conditions, and competes with the veridical depth percept in a classical bistable paradigm that was found to exhibit a high degree of hysteresis. Under the illusory depth percept, reverspectives appear to move vividly as the viewer moves in front of them. This paper reports two experiments that were designed to assess the effectiveness of the painted cues in eliciting the illusory depth percept by using three different measures for the strength of the illusion. As expected, the illusion was favored by monocular viewing and large viewing distances. The results from these two experiments are in close agreement with each other, and they indicate that the painted cues are powerful in influencing the ultimate percept.

1 Introduction

Patrick Hughes, the English artist, has devised an ingenious class of paintings, that are best known as *reverspectives* (Slyce 1998; Wade and Hughes 1999). The simplest way to introduce this class is to describe his first such painting, created in 1964 under the title *Sticking Out Room* (Slyce 1998, page 33). Rather than being painted on a flat surface the room is depicted on a piecewise-planar surface: a truncated pyramid with its large square base on the plane of the wall where the artwork is hung, and the small square face closer to the viewer than the base. Both the base and the small face are frontoparallel surfaces when the viewer stands right in front of the artwork. Hughes painted the small face as the far interior wall of a room. He also painted the two vertical planes that connect the pairs of vertical edges of the base and the small surface as the interior walls adjacent to the far wall. The picture is completed by painting the top and bottom slanted planes as the room's ceiling and floor, respectively.

Hughes's ingenuity is evident when we consider the competition between the actual three-dimensional (3-D) structures of the artpiece on one hand, and the depicted scene, as portrayed by the perspective information painted on the surfaces, on the other hand. The competition results from the opposite 3-D percepts that these cues elicit in the viewer's mind. The 3-D construction of the artpiece, by itself, elicits the veridical percept of a truncated solid convex pyramid that protrudes out of the wall. Stereoscopic viewing, mainly, as well as other cues, such as motion parallax, size, accommodation, and convergence, enhance this veridical percept, in which the small face is nearer than the base. On the contrary, the painted room, by itself, elicits the false illusory percept of a rectangular hollow concave space that recedes into the wall. The strong perspective cues enhance this illusory percept, in which the small face is further away than the base.

In perspective drawing, a vertically oriented rectangle that does not have a frontoparallel orientation in the real world appears as a trapezoid with its far vertical edge depicted as smaller than the near vertical edge; ditto for a horizontally oriented rectangle.

The slanted surfaces of the *Sticking Out Room* (side walls, ceiling, and floor) produce retinal images of trapezoids, just as rectangles would under perspective. However, the edge of the trapezoid which has a retinal image shorter than that of the opposite edge is physically *closer* than that opposite edge. Thus the perspective cues compete with the actual 3-D piecewise-planar surface of the artpiece, hence the name 'reverspectives'.

The most fascinating aspect of reverspectives is that, under the illusory depth percept, they appear to *move vividly* as the viewer moves in front of them. This illusory motion results from the false depth percept and the ambiguity of head-movement parallax (Gregory 1970; Wallach 1976, 1987; Rock 1983; Gogel 1990; Wade and Hughes 1999; Papathomas 2000; Wexler et al 2001). Head-movement parallax, or simply parallax, is the relative movement of the retinal images of features in the outside world that results from the observer's head movement. A given parallax may result from several possible stimuli. Usually, only the veridical stimulus is perceived by the observer. However, there are situations where a given parallax elicits other stable percepts, in addition to the veridical one. This is expected to be the case if the depth percept is not veridical. Reverspectives offer a prime example of a stimulus that yields a false depth percept under a wide range of viewing conditions, such as monocular and binocular viewing, viewing distance and angle, etc. If viewers maintain the false depth percept as they move in front of the reverspective, then they perceive the scenery as moving in a manner that is consistent with the parallax signal. It is interesting that the combination of the false depth percept and the parallax signal overrides the rigidity assumption that plays such as important role in the integration of local motion signals into the percept of global motion (Weiss and Adelson 2000).

The objective of this paper is to examine the role of the painted cues in eliciting the false depth percept. This role was quantified in two separate experiments that used three different measures for assessing the strength of the illusion. The results from both experiments provide strong evidence for a significant role for the painted cues.

2 Stimuli

Both experiments were conducted with Patrick Hughes's *Beyond the Edge* (1999; see http://www.perceptionweb.com/perc0999/wade.pdf), which is shown in figure 1. Experiment 2 also used the stimuli shown in figure 2. The stimulus was hung on a wall, indicated by the hatched pattern in figure 1. The pleated surface on which doors alternate with partly cloudy skies was fastened on the white backplane along the long vertical edges. The pleated surface was constructed from an appropriately shaped planar pattern that was created at the indicated vertical edges. The backplane measured 20.2 cm by 24.0 cm. The pleated surface forms hollow spaces between itself and the backplane. The rest of the dimensions are shown in the figure, which is not drawn to scale. The 14.0 cm by 17.4 cm rectangle defined by the four extreme vertices of the pleated pattern was painted with the same color as the sky.

A casual glance at the front view of *Beyond the Edge* in figure 1 will probably elicit the illusory percept of the short vertical edges, the 'hinge edges', portrayed behind the long ones, the 'door-knob edges'. This is a classical case of a reverspective, in which the near edges are physically shorter than the far edges, and the depicted perspective cues conspire to produce the false depth percept. To study the role of the painted cues in the painting, I used three versions of stimuli that have exactly the same geometrical structure as figure 1. What sets these versions apart is the degree to which they contain painted cues, as explained in what follows (version 1 was painted, whereas versions 2 and 3 were constructed with white paper). (1) The 'fully painted' version is the one sketched in figure 1b. Doors were painted yellow, set against blue skies and grey-and-white clouds. (2) The 'edge-only' version, which lacked all the color and painted features shown in figure 1b, except for the seven vertical edges of figure 1b (four long and three

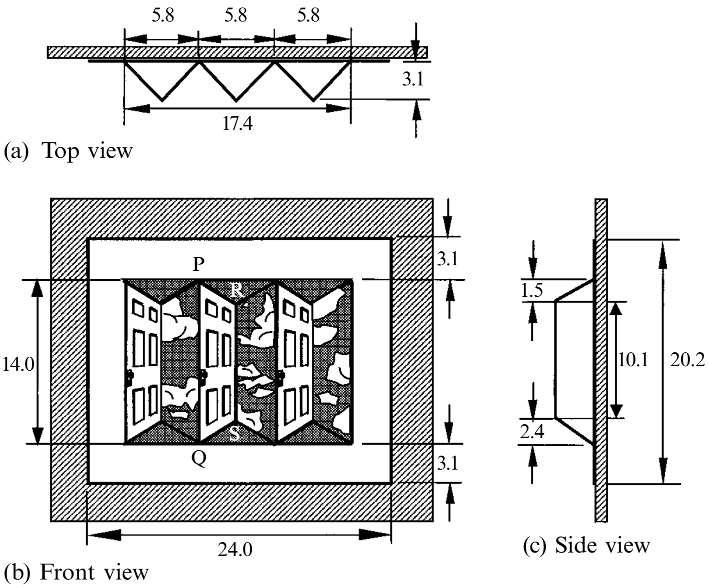

Figure 1. Orthographic views for the reverspective *Beyond the Edge*. Top view (a), front view (b), and side view (c). All dimensions are in cm. Vertical edge RS is physically closer to the observer than PQ. Dimensions indicated on the drawing are accurate, but the figure is not drawn to scale. The wall on which the artpiece is hung is indicated by the hatched pattern. 'Depth' = distance from near vertical edges (such as RS) to backplane = 3.1 cm.

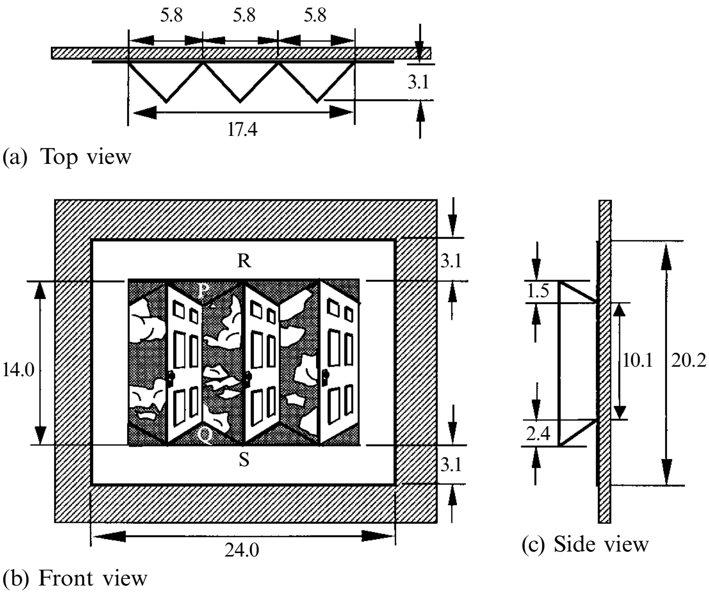

Figure 2. Orthographic views for the properspective version of *Beyond the Edge*. Top view (a), front view (b), and side view (c). Vertical edge RS is physically closer than PQ. The same conventions are used as in figure 1.

short ones) and the zigzag top and bottom edges of the pleated surface, for a total of nine edges. These nine edges were accented by thick black lines on uniformly white surfaces. (3) The 'blank' version, in which the painted cues were eliminated altogether.

In addition to reverspectives, another class of 3-D objects was employed in experiment 2, by simply folding the surfaces in the opposite direction (see figure 2).

In this class, the perspective cues are proper, ie consistent with the 3-D structure of the object. In the spirit of Hughes's notation, these objects will be referred to as 'properspectives'. Figure 2 shows the three main orthographic projections of the fully painted properspective version of *Beyond the Edge*. With the 3-D structure of figure 2, I used the analogs of the three versions that were used with reverspectives: blank, edge-only, and fully painted versions.

3 Experiment 1

This experiment was conducted with four observers who had extensive exposure to the stimuli prior to data collection. Three observers were naïve as to the purposes of the experiment, the fourth was the author. They all had normal, or corrected-to-normal, acuity and possessed functional stereopsis. 'Critical distance' was used as a measure of the strength of the illusion (Hill and Bruce 1993, 1994; Papathomas and DeCarlo 1999). This is the viewing distance at which the percept switches between veridical and illusory. If the illusion is strong, then it will be maintained even when observers get quite close to the reverspective; thus, the stronger the illusion, the smaller the critical distance.

3.1 *Methods*

Each version (fully painted, edge-only, or blank) was placed on a wall at eye level. No fixation point was used. Observers viewed the piece always from a frontal view, ie the piece was straight ahead of them, for all viewing distances. Two viewing conditions were used: monocular and binocular. Observers were asked to report a switch in percept as soon as it occurred, under two viewing schemes. (1) *Retreat*. In each retreat trial, they started from a very short distance (5 cm) away from the piece, at which they all had the veridical percept. Then they started moving away at a slow pace, until the percept switched to the illusory one. The distance D_r at which this switch occurred was recorded, and the next trial commenced. (2) *Approach*. In this type of trial, observers started from a long distance away from the piece (at least 20 cm further away than D_r), at which we ensured they had the illusory percept. Then they started walking towards the piece at a slow pace, until the percept switched to the veridical one; at that moment we recorded the distance D_a. Distances were measured from the observer's eyes to the nearest point of the artpiece, with a measuring tape that was glued on the floor. To ensure that they were certain of their percept at all times, observers were asked to move their heads left and right with an amplitude of 5 to 10 cm. This motion would produce an illusory motion of the piece itself, provided that they had the false depth percept, and this illusory motion was much easier to report than illusory depth.

Thus the experimental design entailed three stimulus types (fully painted, edge-only, blank) × two viewing conditions (monocular, binocular) × two viewing schemes (retreat, approach), for a total of twelve conditions. There were eight trials for each condition, and the order of presentation for stimulus type and viewing mode was randomized across trials. As far as viewing scheme was concerned, trials were run in pairs of successive approach and retreat schemes for the same stimulus type and viewing condition, and the order of approach and retreat was randomized across trial pairs.

3.2 *Results*

Results for all four observers are shown in the four panels of figure 3. The graphs plot the critical distance as a function of the experimental condition. Each point on the graph represents the average of the critical distances for eight trials. Connecting the graph points is unconventional, since the horizontal axis does not correspond to an interval scale. However, the connecting lines highlight the pattern of dependency of the critical distance on the experimental conditions. Some points appear to be missing altogether from the graphs; actually, they can be thought of as approaching 'infinity'.

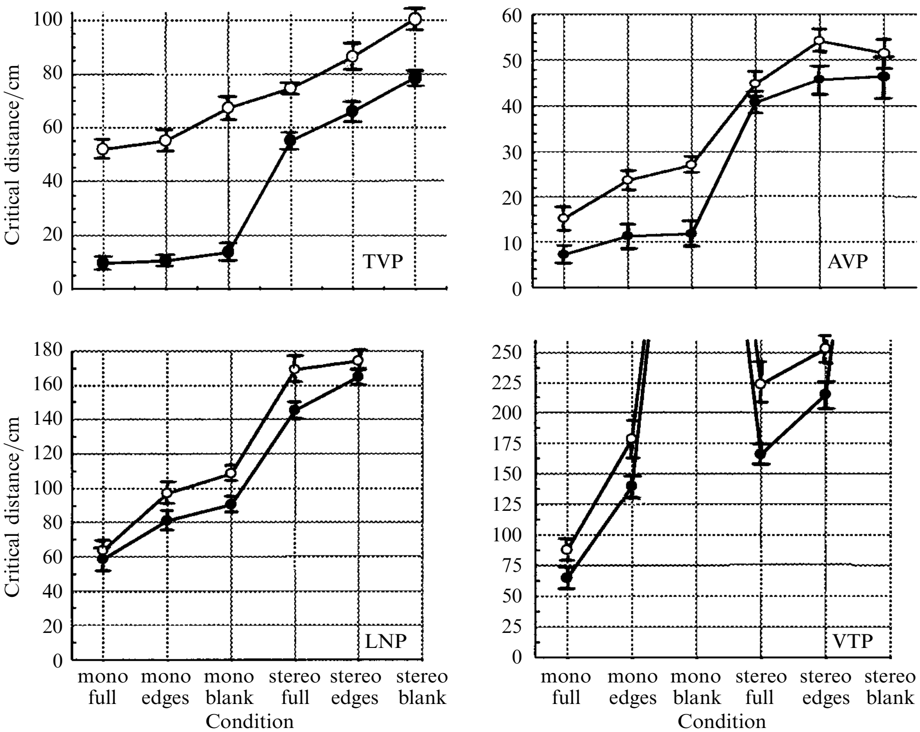

Figure 3. Results of experiment 1 for four observers. Graphs plot the critical distance for the twelve conditions: three stimulus types ('full'—fully painted; 'edges'—edge-only painted; 'blank'—unpainted) × two viewing conditions ('Mono'—monocular, left group of three; 'Stereo'—binocular, right group of three) × two viewing schemes (retreat—open symbols; approach—solid symbols). Error bars represent one standard error of the mean.

This means that the distance for obtaining the illusory percept was 'infinite', ie the stimulus was always perceived veridically, no matter how far away the observer was (the maximum possible viewing distance in the experiment was 9.5 m).

The importance of the painted cues is evident from the graphs. An analysis of variance (ANOVA) test was conducted to assess the role of painted cues for each of the four conditions: monocular approach (MA), monocular retreat (MR), binocular approach (BA), and binocular retreat (BR). These tests confirmed that the critical distance decreases significantly, ie the illusion gets stronger, as one goes from the blank to the edges-only to the fully painted stimuli. Painted cues were highly significant for observers LNP and VTP ($p < 0.0005$ in all four conditions). They were also significant for observer TVP ($p < 0.01$ for three conditions, and statistical significance was only just missed for condition MA, for which $p < 0.051$). For observer AVP the effect of painted cues on critical distance was statistically significant for conditions MR ($p < 0.01$) and BR ($p < 0.018$), but not for MA ($p < 0.101$) or BA ($p < 0.095$). The role of the painted cues is, remarkably, stronger for the observers who indicated 'infinite' critical distances for the blank stimulus: for observer LNP this occurred under binocular viewing, and for VTP under both monocular and binocular viewing. For these two observers, the painted cues, even in the form of edges in the edge-only stimuli, are *necessary* for obtaining the illusion.

3.3 *Discussion*

The trends in the data are largely as one would expect. In particular, critical distances are significantly smaller under monocular than binocular viewing. The fact that the critical distance for the approach scheme is always smaller than that for the retreat scheme reveals hysteresis in the percept, which is typical of bistable stimuli (Fender and Julesz 1967; Ditzinger and Haken 1989; Hill and Bruce 1993; Papathomas and DeCarlo 1999). Observers start with the illusory percept at large distances in the approach scheme, and they maintain the illusion up to a distance D_a, at which the switching occurs. On the other hand, in the retreat scheme they begin with the veridical percept because they start from a very small distance, and they maintain the veridical percept as they retreat up to a distance D_r, where $D_r > D_a$.

The data obtained from this experiment demonstrate the need for different measures of the strength of the illusion. The main problem has to do with the substantial interobserver differences in performance. The absence of a fixation point may have increased the variability in the data. In addition, 'infinite' distances were obtained for some conditions for two observers, which did not allow results to be averaged across observers. These problems were addressed in the design of experiment 2.

4 Experiment 2

This experiment was designed to obtain data by using two additional measures for the strength of the illusion: first, the initial percept obtained immediately after the observer was exposed to the stimulus; and, second, the relative 'residence time' (Kruse et al 1994; Kovács et al 1996; Logothetis et al 1996) of the illusory percept as a percentage of the total viewing time, as they alternated stochastically in time.

4.1 *Methods*

In addition to assessing the role of painted cues in 3-D reverspective objects, this experiment allowed me to assess their role in properspectives as well (see section 2 and figure 2). In these objects, the perspective cues are in agreement with the actual 3-D structure of the surface, hence they *enhance* the veridical depth percept, rather than *compete* against it, as they did in reverspectives. Thus, in experiment 2 I used two perspective formations (reverse and proper) and for each formation I used three painting versions (fully painted, edge-only, and blank) for a total of six types of objects.

A fixation cross was drawn at the same relative location on the three reverspectives near the center of the middle door, and on the three properspectives near the center of their middle door. Two small square colored patches, one red and the other green, 0.4 cm on each edge, were attached 1.1 cm to the right and to the left of the fixation cross. It is obvious that the perceived global depth percept (veridical or illusory) is perfectly correlated with the perceived depth ordering of the colored patches. This was verified empirically with numerous pilot trials, during which observers reported simultaneously both the global percept and the local depth arrangement of the two patches.

Each object was placed on a wall at the same height as the observer's eyes. Observers sat on a chair in front of the object at a distance of 127 cm for monocular viewing, and 254 cm for binocular viewing. To avoid shadows being cast on the objects, observers wore a head-mounted flashlight, in a near-dark room. Observers started with their eye(s) closed while the experimenter placed each object on the wall. When the object was set, they were instructed to open their eye(s), fixate on the fixation point, and commence the trial. They viewed each object for at least one minute, during which they were told to press two different buttons, depending on which color patch appeared in front. Thus, a series of instants at which the two buttons were pressed was obtained, indicating the starting points and the durations of each percept. Depending on the observer, some stimuli resulted in several button presses, indicating a strongly bistable percept. Other stimuli had a stable dominant percept and did not exhibit any alternations. The trial ended either immediately upon the first button press (stimulus switch) that occurred after the 60 s mark, or at 90 s, whichever came first.

Twelve observers took part in this experiment. They were all tested for normal acuity, stereoscopic vision, and color vision [for the latest they had to recognize all fifteen numerals in pseudo-isochromatic (AOC 1965) Ishihara plates]. The experiment was conducted in two groups of Latin-square designs with six observers × six types of stimuli. There were two trials for each of the six types of objects under both monocular and binocular viewing. In the first of these two trials, the color of the patch in front was chosen at random and determined which of the two copies of the stimulus was to be used; the other copy was used in the second trial.

4.2 *Results*

The data from experiment 2 were analyzed to produce two measures for the relative strengths of the illusory and veridical depth percepts. (1) The original percept obtained immediately after observers were exposed to the stimulus, as they opened their eyes (figure 4a); the first percept in bistable stimuli is usually the dominant percept (Gale and Findlay 1983; Kanizsa and Luccio 1994). (2) The percentage of the residence time of the veridical percept, obtained as the ratio of the total time that the observer spent in the veridical percept to the total viewing time (figure 4b). By choosing not to terminate the trial at precisely 60 s the last percept was favored; however, this imbalance was equalized when averaged across trials and observers. The advantage of these measures over the critical distance is that there are no undefined values, hence they can be averaged across observers.

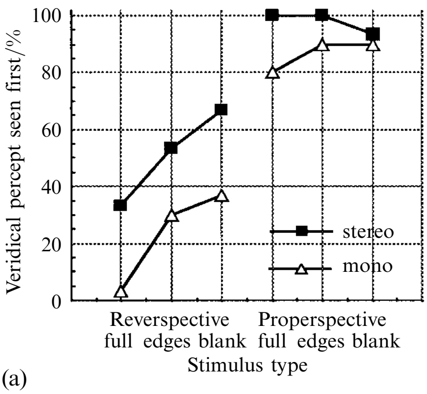

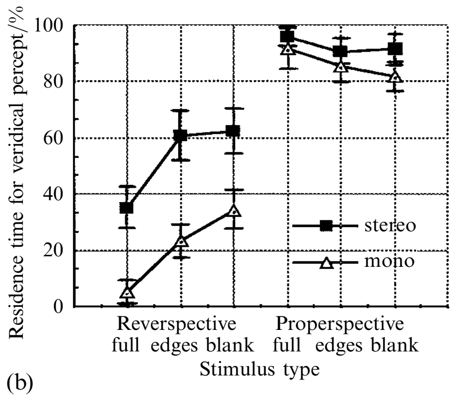

Figure 4. The results of experiment 2. (a) The percentage of trials in which the veridical depth was perceived first. There are not enough data per point to estimate standard errors. (b) The percentage of the residence time of the veridical percept. Triangles and squares are used for monocular and binocular viewing, respectively. Stimulus type is shown along the horizontal axis. Error bars represent one standard error of the mean.

The results for measures 1 and 2 are shown in figures 4a and 4b, respectively, averaged across trials and observers. As with figure 2, the smaller the value on the ordinate axis, the stronger the illusory percept. The important role of painted cues is obvious in both graphs, especially for the objects with the reverspective formation. As expected, the dependency of the strength of the illusion on the stimulus follows the same trends in the graphs of figures 4a and 4b, reinforcing the validity of using the two measures. Standard errors are not shown in figure 4a because there are only twenty-four data per point (two trials for each of twelve observers).

4.3 *Discussion*

The main effect of painted cues is evident from a comparison of the data between the fully painted and the blank versions in the same perspective formation. Since the painted cues are mainly of perspective nature, the effect of these cues is different for the two formations. In reverspectives, painted cues compete against the veridical percept, hence they enhance the illusion; in properspectives, on the contrary, they are in agreement with the true 3-D geometry and hence they enhance the veridical percept. These trends are indeed evident in the graphs of figure 4. However, the role of the painted cues is much stronger in enhancing the illusion in reverspectives than in enhancing the veridical percept in properspectives. In the former, the increase in the mean residence time of the illusion between the blank and fully painted versions is from 37.8% to 65.2% for binocular and from 65.9% to 95.0% for monocular viewing, ie an increase of 27.4 and 29.1 percentage units, respectively. The corresponding figures in properspectives are an increase in the mean veridical residence time from 91.5% to 95.7% (binocular) and from 81.6% to 91.5% (monocular), ie 4.2 and 9.9 percentage units, respectively. These differences arise most likely from a saturation effect: in properspectives, the veridical residence times are already high because the percept for the nonpainted surface is already strongly veridical (better than 81.6%); thus the painted cues have little space to make substantial improvements. On the contrary, the veridical residence times are quite low in the reverspective case (less than 65.9%), giving the painted cues much latitude to exhibit their effect.

This is reflected in the results of the statistical analysis. A repeated-measures ANOVA was carried out on the data of figure 4b that pertain to reverspectives. The effect of the painted cues, obtained by comparing the fully painted, edges-only, and blank conditions, was found to be statistically significant for both monocular ($F = 11.039$, $p < 0.0005$) and binocular viewing ($F = 4.071$, $p < 0.031$). The corresponding test on the data obtained with properspectives failed to show statistical significance ($F = 2.092$, $p < 0.147$ for monocular viewing; $F = 2.062$, $p < 0.151$ for binocular viewing).

As in experiment 1, the data differences between binocular and monocular viewing are exactly as expected, with binocular viewing always favoring the veridical percept. This difference is larger for the reverspectives than the properspectives, and it is probably due to a saturation effect, similar to the one presented in the above paragraph.

If one uses the original-percept measure when experimenting with a small number of observers and trials, the data are not adequate to assess the relative strengths of the two percepts. The main problem is that the initial percept may depend on accidental factors or on priming (Ditzinger and Haken 1989). In such cases, the percent-residence-time index is a much more reliable measure (Kovács et al 1996; Logothetis et al 1996). In the present case, with twelve observers and two trials per condition, the dependency of the strength of the illusion on the stimulus generally follows the same trends in the graphs of figure 4a and 4b, indicating that the original-percept measure provides a relatively reliable index. Nevertheless, the data of figure 4b for percent residence time are more representative of the relative strengths of the two percepts, as they reflect some measure of the dynamic alterations that depend precisely on these relative strengths.

5 Conclusions and discussion

Results from the two experiments are well correlated in terms of the dependency of the strength of the illusion on the viewing conditions and the object used as stimulus. The three measures (critical distance, original percept, and percent residence time) produced comparable results, and confirmed the important role of the painted cues. For all three measures, the painted cues are seen to have statistical significance by comparing the results between the fully painted and the unpainted versions of the reverspective. Even though there is a clear increase in the mean value of the percent residence time obtained with the fully painted version over the unpainted version of the properspective, this increase did not reach statistical significance, probably for the reason presented in section 4.3.

An interesting issue is what painted cues, other than foreshortening, are important in *Beyond the Edge*, and what their relative strength is in enhancing the false depth percept.

In addition to the foreshortening cue that is also present in the edge-only version, the fully painted *Beyond the Edge* has (1) more perspective cues, ie the horizontal lines of the door panels that meet at the vanishing point, (2) recognizable objects (doors), (3) the juxtaposition of doors and skies, (4) the shadows on the doors, and (5) the natural colors of the painting, among others. Apparently, some of these cues, and/or their combination, make the difference in performance between fully painted and unpainted reverspectives significant. In fact, *Beyond the Edge* is relatively poor, as reverspectives go, in terms of painted cues. Most reverspectives have shadows that are consistent with the illusory percept, and hence inconsistent with the veridical percept (Slyce 1998; Wade and Hughes 1999; Papathomas 2000). Some contain texture gradients (eg in the form of a receding tiled floor) that provide additional cues for both perspective and foreshortening. Assessing the role of these cues is the subject of experiments in progress that use much more complex reverspectives.

Along the same lines, an interesting project would be to try to isolate the roles of other low-level signals, such as disparity and head-motion parallax, in addition to the painted signals. An integrated approach would be to study systematically the interaction of painted cues and the other low-level signals, to extend the results of earlier studies such as those by Young et al (1993), Johnston et al (1994), Landy et al (1995), and Sperling and Dosher (1995). At the other extreme, experiments can be designed to study whether cognitive factors, such as familiarity with objects and scenes, prior experience, priming, memory, etc, influence the percept. From a Bayesian point of view, the role of prior probabilities, such as the bias to interpret intercepting lines as orthogonal and parallelograms as rectangles (Griffiths and Zaidi 1998, 2000) may also play a role in shaping the final percept.

Both experiments revealed significant interobserver differences with respect to the relative strength of the illusory and the veridical percepts in reverspectives. The former appears to be favored by painted signals such as perspective, foreshortening, texture gradient, and shadows. The latter is favored by the stereoscopic disparity signals. The head-motion parallax signals that were present in experiment 1 (but not in experiment 2) are ambiguous, because they are consistent with either of the two percepts. However, these parallax signals, which are identical as far as the retinal flow fields are concerned, are interpreted entirely differently by the visual system, depending on what is the prevalent depth percept. Under the veridical depth percept, the motion signals are processed to produce the percept of a stable, stationary, rigid object. On the other hand, under the illusory depth percept, the same motion signals produce the percept of a moving, nonrigid object. The object's surfaces undergo strange distortions in this case, because of the false depth percept. This is akin to the illusory motion of hollow face masks (Gregory 1970; Hill and Bruce 1993, 1994; Papathomas and DeCarlo 1999) or 3-D wire-frame Necker cubes (Rock 1983). Thus, if the visual system favors rigidity and stable surfaces, one may argue that the motion parallax signal favors the veridical percept. On the basis of the above, it would seem that observers who favor the veridical or illusory percept assign a large or small weight, respectively, to their disparity-driven mechanisms, but this needs to be tested. What other mechanisms contribute to the relative dominance of these percepts? It would be interesting to correlate observers' performance when viewing reverspectives to other performance indices with low-level testing stimuli.

Acknowledgments. I would like to express my gratitude to Patrick Hughes for sharing his ideas on art, for coming to Rutgers University with Dr Diana Atkinson to present a lecture, and for generously supplying me with his artpieces. I am indebted to Lisa Bono for helping to construct the stimuli, planning and administering experiment 2, processing and plotting the data, and helping to edit the manuscript. Many thanks to Nicholas Wade for bringing me in touch with Patrick and encouraging me to work on these projects, and to Akos Feher for technical support.

References

AOC, 1965 *Pseudo-isochromatic Plates for Testing Color Perception 1965* (Philadelphia, PA: American Optical Corporation/Beck Engraving Company)

Ditzinger T, Haken H, 1989 "Oscillations in the perception of ambiguous patterns: A model based on synergetics" *Biological Cybernetics* **61** 279 – 287

Fender D, Julesz B, 1967 "Extension of Panum's fusional area in binocularly stabilized vision" *Journal of the Optical Society of America* **57** 819 – 830

Gale A G, Findlay J M, 1983 "Eye movement patterns in viewing ambiguous figures", in *Eye Movements and Psychological Functions: International Views* Eds R Groner, C Menz, D F Fisher, R A Monty (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 145 – 168

Gogel W C, 1990 "A theory of phenomenal geometry and its applications" *Perception & Psychophysics* **48** 105 – 123

Gregory R L, 1970 *The Intelligent Eye* (London: Weidenfeld & Nicolson)

Griffiths A F, Zaidi Q, 1998 "Rigid objects that appear to bend" *Perception* **27** 799 – 802

Griffiths A F, Zaidi Q, 2000 "Perceptual assumptions and projective distortions in a three-dimensional shape illusion" *Perception* **29** 171 – 200

Hill H, Bruce V, 1993 "Independent effects of lighting, orientation, and stereopsis on the hollow-face illusion" *Perception* **22** 887 – 897

Hill H, Bruce V, 1994 "A comparison between the hollow-face and 'hollow-potato' illusions" *Perception* **23** 1335 – 1337

Johnston E B, Cumming B G, Landy M S, 1994 "Integration of stereopsis and motion shape cues" *Vision Research* **34** 2259 – 2275

Kanizsa G, Luccio R, 1994 "Multistability as a research tool in experimental phenomenology", in *Ambiguity in Mind and Nature: Multistable Cognitive Phenomena* Eds P Kruse, M Stadler (Berlin: Springer) pp 47 – 68

Kovács I, Papathomas T V, Yang M, Feher A, 1996 "When the brain changes its mind: Interocular grouping during binocular rivalry" *Proceedings of the National Academy of Sciences of the USA* **93** 15508 – 15511

Kruse P, Struber D, Stabler M, 1994 "The significance of perceptual multistability for research on cognitive self-organization", in *Ambiguity in Mind and Nature: Multistable Cognitive Phenomena* Eds P Kruse, M Stadler (Berlin: Springer) pp 69 – 84

Landy M S, Maloney L T, Johnston E B, Young M J, 1995 "Measurement and modeling of depth cue combination: In defense of weak fusion" *Vision Research* **35** 389 – 412

Logothetis N K, Leopold D A, Sheinberg D L, 1996 "What is rivalling during binocular rivalry?" *Nature* **380** 621 – 624

Papathomas T V, 2000 "See how they turn: False depth and motion in Hughes's reverspectives" *Human Vision and Electronic Imaging V* SPIE Proceedings Series, volume **3959** 506 – 517

Papathomas T V, DeCarlo D, 1999 "Top – down and bottom – up processes in 3-D face perception: psychophysics and computational model" *Perception* **28,** supplement 112

Rock I, 1983 *The Logic of Perception* (Cambridge, MA: MIT Press) pp 9 – 11

Slyce J, 1998 *Patrick Hughes: Perspective* (London: Momentum)

Sperling G, Dosher B A, 1995 "Depth from motion", in *Early Vision and Beyond* Eds T V Papathomas, C Chubb, A Gorea, E Kowler (Cambridge, MA: MIT Press) pp 133 – 143

Wade N J, Hughes P, 1999 "Fooling the eyes: *trompe l'oeil* and reverse perspective" *Perception* **28** 1115 – 1119

Wallach H, 1976 "The apparent rotation of pictorial scenes", in *Vision and Artifact* Ed. M Henle (New York: Springer) pp 65 – 69

Wallach H, 1987 "Perceiving a stable environment when one moves" *Annual Review of Psychology* **38** 1 – 27

Weiss Y, Adelson E H, 2000 "Adventures with gelatinous ellipses—constraints on models of human motion analysis" *Perception* **29** 543 – 566

Wexler M, Panerai F, Lamouret I, Droulez J, 2001 "Self-motion and the perception of stationary objects" *Nature* **409** 85 – 88

Young M J, Landy M S, Maloney L T, 1993 "A perturbation analysis of depth perception from combinations of texture and motion cues" *Vision Research* **33** 2685 – 2696