Perception of Surfaces from Line Drawings

CHRISTOPH HOFFMANN[1], ZYGMUNT PIZLO[2],
VOICU POPESCU[1], STEVE PRICE[1]
[1]Computer Sciences, [2]Psychological Sciences, Purdue University

We test the perception of 3D surfaces that have been rendered by a set of lines drawn on the surface. Each surface is rendered as a family of curves which are in the simplest case the intersections with a family of parallel planes. On each trial, a surface or its "distorted" version is shown in this way, in an arbitrary orientation on an LCD screen or in a volumetric 3D display. The distortion is produced by stretching the surface in the z-direction by 30%. The subject's task is to decide whether two sequentially presented surfaces are identical or not. The subject's performance is measured by the discriminability *d'*, which is a conventional dependent variable in signal detection experiments. The work investigates the question whether a surface rendered with planar and geodesic curves is easier to recognize than one where the curves are not planar or not geodesic.

General Terms: psychophysical experiment, shape perception, line drawing, volumetric display, 3D image.

1. INTRODUCTION

This paper addresses a question of perceptual reconstruction of 3D surfaces. The reconstruction problem is computationally difficult because the 3D percept has to be produced from 2D image(s). It is known that this inverse problem can be solved (at least in principle) if the visual system can impose constraints on the family of possible solutions (see Pizlo, 2001 for a review). To shed more light on the underlying perceptual mechanisms we study the effect of constraints that can be applied to surface contours: planarity and geodesic constraints. We also test the role of binocular disparity as a depth cue. Binocular viewing is tested by using Perspecta, a volumetric display.

The paper is organized as follows: Section 2 reviews prior work, Section 3 describes the psychophysical experiment conducted, Section 4 presents results, Section 5 provides discussion, and Section 6 sketches possible directions for future work.

This research was supported in part by NSF grants DMS-0138098, EEC-0227828, EIA-0216131, and ACI-0325227. Hoffmann is also supported in part by an IBM Faculty Award.
Authors' addresses: Christoph Hoffman, Zygmunt Pizlo, Voicu Popescu, Steve Price, 250 N. University Street, West-Lafayette, IN, 47907.

2. PRIOR WORK

The systematic study of the role of surface contours in perception of 3D surfaces started with the work of Stevens (1981, 1986). He discussed the role of planarity and geodesic constraints, especially in the case of developable surfaces. The effect of geodesic constraint was further studied by Knill (1992, 2001). The interaction of *a priori* constraints imposed on surface contours and binocular disparity was tested by Stevens & Brookes (1988), by Mitchison (1988) and by Pizlo, Li and Franics (2005). Finally, the role of symmetry of an object and its contours was studied by Hochberg & McAlister (1953), Attneave & Frost (1969), and Pizlo, Li & Chan (2005). All these studies demonstrated that contours constraints are critical not only in monocular, but also in binocular vision.

3. PSYCHOPHSYSICAL EXPERIMENT

3.1 Subjects

Five subjects were tested including one author (SP). SP was familiar with the stimuli and with the research hypotheses being tested. The other four subject were naïve as to the design of stimuli and the hypotheses. SP received substantially more practice than the other four subjects.

3.2. Stimuli

The surfaces to be rendered are a family of single Gaussian functions with different aspect ratios. Given a Gaussian function F that is restricted to a standard domain, the intersection with a family of intersecting surfaces is computed. In the simplest case, the intersecting surfaces are a family of parallel planes, but in more complex cases other surface families are used. The number of intersecting surfaces was constant, but their position relative to the Gaussian surface, as well as orientation relative to the square base was randomized, in order to avoid comparing local cues, rather than the shapes of the whole surfaces. The intersection with a particular plane is computed using a simplicial continuation method; see, e.g., Allgower and Gnutzmann (1991). The method is related to the well-known "marching cubes" method from computer graphics, e.g., Bloomenthal (1994), but by subdividing into simplices the ambiguous cases are avoided.

The implementation of the continuation method assumes only that the manifold of simplices is topologically a disk. This is easily accomplished in the case of planar sections. We extended it to nonplanar surfaces by simplicial subdivisions of annular regions which were cut to be topologically a disk. The seam along which the annulus was cut requires no special treatment as long as the discretization along the seam is compatible. That is, the fact that an intersection curve crossing the seam is connected can be ignored by the rendering algorithm and the result is indistinguishable by the observer.

Five families of intersecting surfaces were considered:

1. Parallel vertical planes that are parallel to the axis of symmetry of the Gaussian *F*.
2. Parallel oblique planes intersecting the symmetry axis of the Gaussian at an angle of 45 degrees.
3. Radial vertical planes that are parallel to the axis of symmetry of the Gaussian *F*.
4. Radial oblique planes intersecting the symmetry axis of the Gaussian at an angle of 45 degrees.
5. A family of spheres.

The family of spheres consists of spheres of equal radius whose centers are along a line and are evenly spaced. The center line lies in the plane z=0 and intersects the axis of symmetry of the Gaussian. Examples for each family of curves are shown in Figure 1.

Contours produced by intersecting surfaces 1-4, but not 5, were planar. All contours in case 3 were geodesic lines. None of the contours in case 2, or 5 were geodesic. One contour in 1 and 4 (the one approximately intersecting the symmetry axis of the Gaussian surface) was approximately a geodesic line.

1	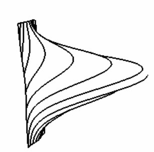	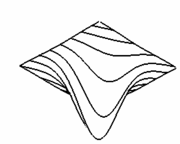
2	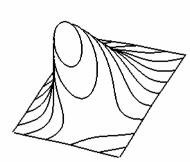	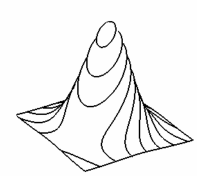

3	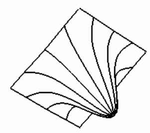	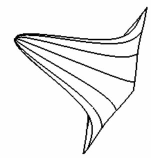
4	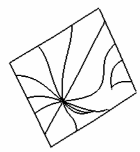	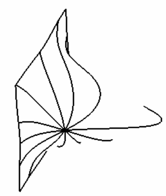
5	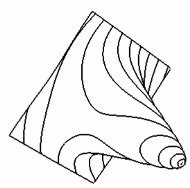	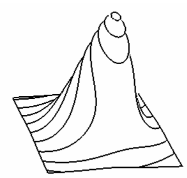

Figure 1 Gaussian function rendered by the intersection lines with a family of parallel vertical planes (1), parallel oblique planes (2), radial vertical planes (3), radial oblique planes (4), and spheres (5).

3.3. Procedure

On each trial the subject was shown two stimuli and the task was to decide whether their aspect ratios were the same. Each stimulus was shown for one second, and they were separated by a one second pause (blank display). The 3D orientation of each stimulus was random subject to some constraints in order to eliminate views that provide zero, or close to zero information about the 3D shape (see Section 3.4). The size of each stimulus was also randomized. As a result, the subject had to pay attention to the aspect ratio of the 3D surface, rather than to its height.

Signal detection method was used (Macmillan & Creelman, 2005). On “same” trials, the two stimuli had identical aspect ratio, and on “different” trials the aspect ratios were different by 30%. The order of trials was randomized. Each session consisted of 200 trials: 100 same and 100 different. Hits and false alarm rates were used to estimate the discriminability d′. Viewing was either monoscopic (binocular viewing of an image displayed on an LCD monitor, see Section 3.5), or stereoscopic (binocular viewing of an image displayed in a volumetric 3D display, see Section 3.6). The order of the 10 sessions (five types of contours and two modes of viewing) was random and different for each of the five subjects.

3.4. View selection

The random views at which an observer sees the rendered Gaussians exclude the case where the planar curves are seen edge-on, with a view direction that lies within *a* degrees of the cutting planes. Such a view would not give any spatial information on account of the intersection curves being a collection of straight-line segments. We also exclude a view that is within *b* degrees of the axis of symmetry, i.e., seeing the Gaussian from above, a view within *c* degrees of being perpendicular to the axis of symmetry, and a view that sees the back face of the Gaussian base plane, since such views again would yield little or no spatial information (Figure 2). In practice, we choose the angle limits *a, b,* and *c* to be 30, 30, and 5 degrees, respectively.

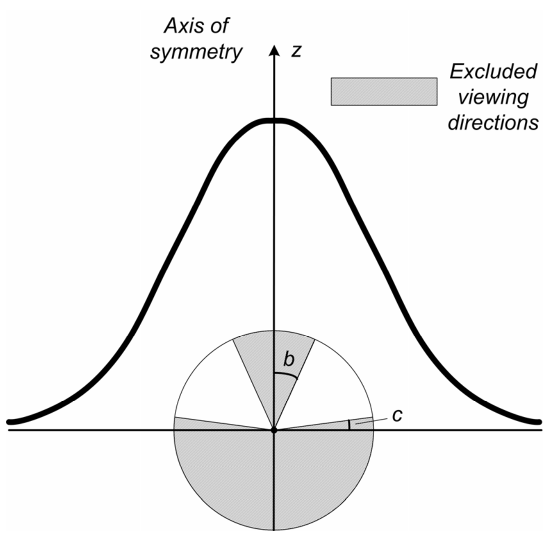

Figure 2 Gaussian viewing direction restrictions.

3.5. Display on conventional LCD

In half of the sessions the images were displayed on a conventional LCD. Viewing was monoscopic, with a binocular stimulus disparity of zero. The viewing distance was approximately 50cm.

Only the visible lines were displayed. The hidden line removal is accomplished by rendering the curves as lines and then rendering the surface itself, using flat shading, in the color of the background, in our case white. The depth ordering and occlusion computations of the graphics hardware then shows the intersection curves only on the visible parts of the surface, irrespective of the point of view. That is, the rendered complex of curves and surface can be freely rotated in real time on standard PC graphics hardware. To resolve numerical issues, the intersections are offset from the underlying Gaussian by a slight dilation. The image is rendered using perspective projection. The projection matrix is computed from the desired viewing distance, desired image resolution (512x512 pixels in our case), and pixel size for the LCD. Figure 1 shows pairs of images displayed during the LCD sessions.

3.6. Display on Perspecta volumetric 3D display

In half of the sessions viewing was stereoscopic. The stimulus was displayed in a volumetric display (Perspecta [Actuality Systems]). The volumetric display builds a 3D image by projecting in rapid succession 2D images on a spinning screen that sweeps the 3D scene. Each 2D image is the intersection of the screen plane with the 3D scene. The inertia of the visual system allows the user to reconstruct the scene which appears like a 3D sculpture of light. Each eye gets the correct image of the scene, without the need of encumbering eyewear and trackers. The viewing distance was approximately 100cm.

The 3D display lacks adequate support for line rendering. For this, the line segments are displayed as thin tessellated cylinders. When two segments share a vertex, the corresponding cylinders share a base face, such that they are connected seamlessly. The cylinders are not shaded (constant color, white in our experiments).

Figure 3 shows photographs of the volumetric display as it renders 3D images of the Gaussian. The screen spins at 24Hz, so we chose the nearest available exposure time which equals 1/25s. During the experiment the lights in the laboratory were dimmed to compensate for the low brightness of the 3D image.

1	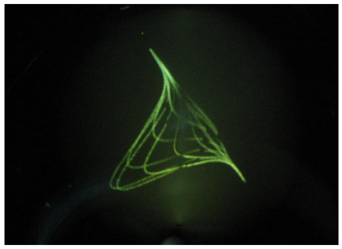	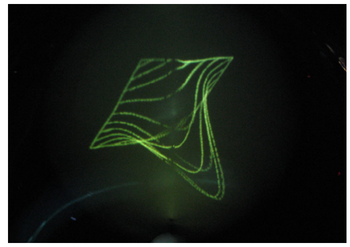
2	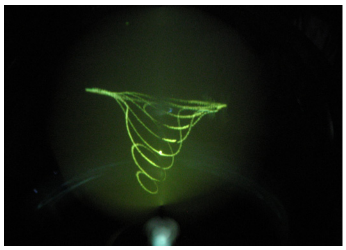	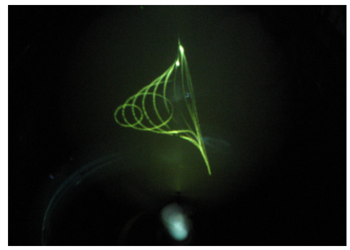
3	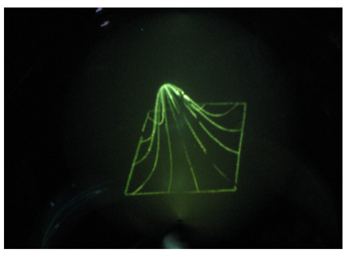	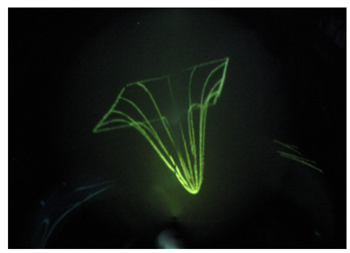
4	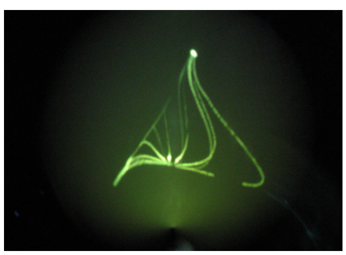	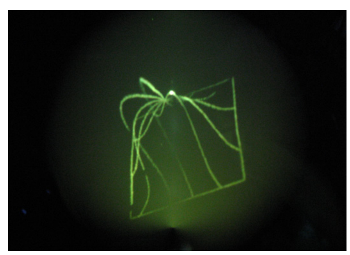
5	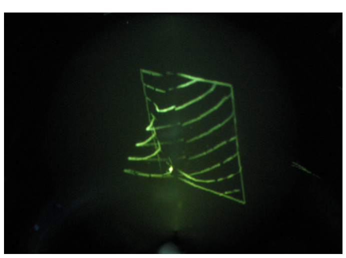	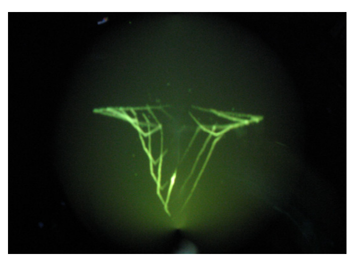

Figure 3 Examples of stimuli rendered on our volumetric display. Each image is a photograph of the display. The Gaussian is rendered by the lines of intersection between the Gaussian and a family of parallel vertical planes (1)), parallel oblique planes (2), radial vertical planes (3), radial oblique planes (4), and spheres (5).

Figure 4 shows a photograph of the experimental setup used for the Perspecta sessions. The 3D image is not visible in the photograph since the laboratory lights were on to better capture the setup. An auxiliary LCD monitor displays the number of the current trial, and the currently selected answer. The subject selects one of the two possible answers *same* or *different* using the left or right mouse buttons. The subject receives auditory feedback for the answer: a high pitched tone indicates that the selected answer was correct, in contrast to a low pitch tone indicative of an incorrect answer.

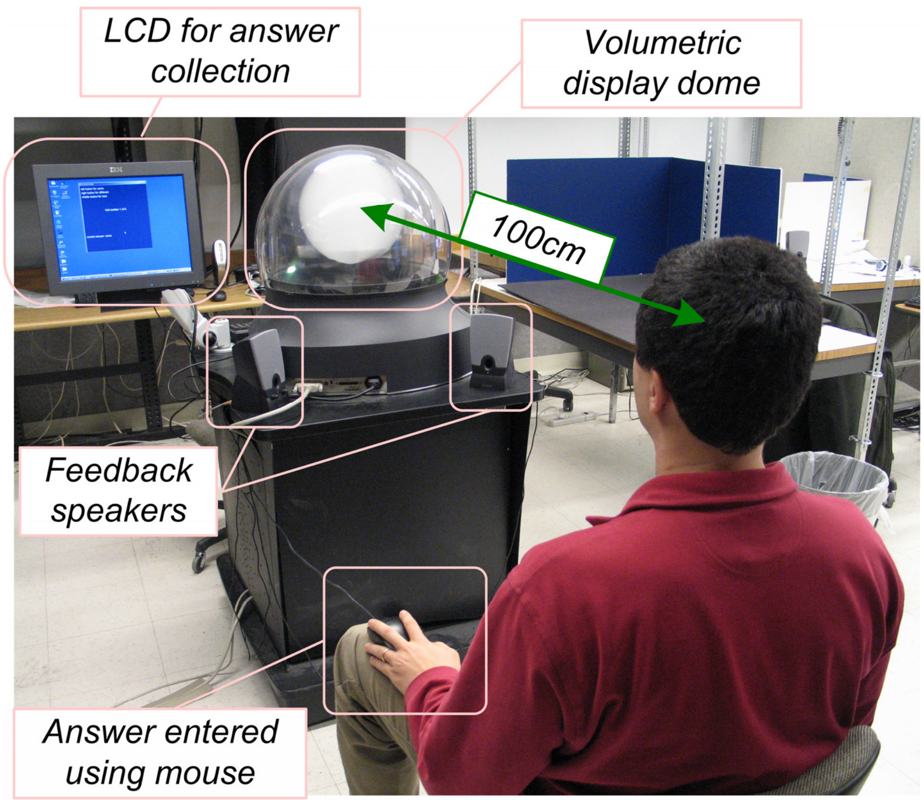

Figure 4 Experimental setup used in the volumetric display sessions.

4. RESULTS

The values of discriminability d′ for each subject and each session are given in Table 1 and Table 2. Average values across all five subjects are plotted in Figure 5.

		Family of intersecting surfaces				
		Parallel vertical planes	**Parallel oblique planes**	**Radial vertical planes**	**Radial oblique planes**	**Spheres**
Subject	**CJ**	0.86	0.41	0.61	0.10	0.65
	CM	1.10	0.57	0.75	0.69	0.78
	JD	1.07	0.31	1.73	0.35	1.51
	SP	1.84	0.78	2.79	0.96	1.92
	WY	0.46	-0.13	0.75	0.40	0.81

Table 1 Discriminability values for LCD (monoscopic viewing), across types of Gaussian visualizations and subjects.

		Family of intersecting surfaces				
		Parallel vertical planes	**Parallel oblique planes**	**Radial vertical planes**	**Radial oblique planes**	**Spheres**
Subject	**CJ**	2.08	0.80	2.82	0.73	2.09
	CM	1.92	1.52	2.57	1.12	1.81
	JD	1.50	0.99	2.03	0.88	0.94
	SP	2.63	1.41	3.44	1.98	3.36
	WY	1.52	0.52	2.08	0.76	0.83

Table 2 Discriminability values for volumetric display (stereoscopic viewing), across types of Gaussian visualizations and subjects.

Examination of Table 1 shows that the pattern of results is similar in all five subjects, although performance of SP is systematically better than that of the other four subjects. Recall that SP was familiar with the stimuli (he actually worked on designing them) and that he received substantially more practice than the other subjects. Having more experience with the stimuli is likely to reduce the variability of the response criterion. When the response criterion is not stable in a single session, the d′ is likely to underestimate the actual discriminability of a given subject. This difference in familiarity led to the overall shift in performance level, without changing the pattern of results.

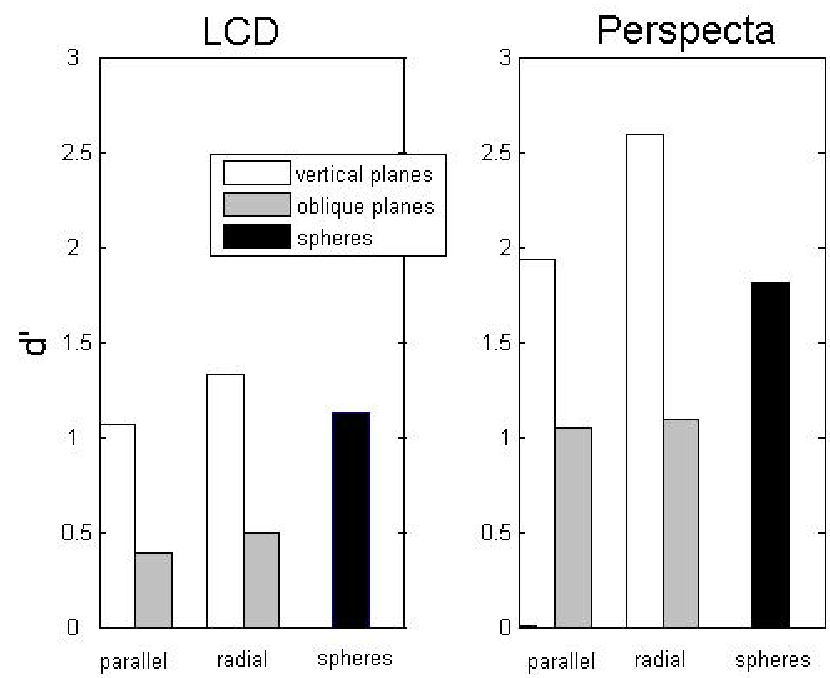

Figure 5 Average performance of five subjects.

Figure 5 shows several effects. To evaluate statistical significance of the effects, we performed two ANOVA analyses. In the first analysis, we applied a 3-factor repeated measures ANOVA to all conditions involving planes. The three factors were: orientation of planes relative to the Gaussian surface (vertical vs. oblique), orientation of planes relative to one another (parallel vs. radial), and viewing mode (monoscopic vs. stereoscopic). The main effect of parallel vs. radial was significant with $p<0.02$ (performance with radial was slightly better), and the main effects of the other two factors were significant with $p<0.001$ (stereoscopic performance was better than the monoscopic one by a factor of about two, and performance with vertical planes was better than that with oblique planes by a factor of two, or more).

Only one interaction was significant ($p<0.05$) and it was the interaction between monoscopic vs. stereoscopic and vertical vs. oblique planes. This interaction means that the difference in performance between vertical and oblique planes depended on whether viewing was monoscopic or stereoscopic. This fact can be seen in Figure 5. In monoscopic viewing, performance with vertical planes was better by 0.5-0.8 than performance with oblique planes. In stereoscopic, this difference was 1.0-1.5. Considering the fact that a total of 7 hypotheses were tested in this analysis, it is reasonable to introduce a correction for type I error. This would mean that effects whose p-value is greater than 0.05/7, should not be considered statistically significant. If this correction is applied, there are only two main effects significant: monoscopic vs. stereoscopic and vertical vs. oblique.

In the second analysis, we applied a 2-factor repeated measures ANOVA. The two factors were: viewing mode (monoscopic vs. stereoscopic) and contour type (vertical parallel, oblique parallel, vertical radial, oblique radial, spheres). The two main effects were significant ($p<0.001$), but the interaction was not ($p=0.29$). This means that the effect of binocular disparity was the same for all five types of contours. Post-hoc tests were performed for all 10 pair-wise comparisons of the 5 types of contours. Eight differences were significant ($p<0.01$). The two non-significant differences were between parallel oblique and radial oblique ($p=0.64$) and between spheres and parallel vertical ($p=0.87$).

5. DISCUSSION

The fact that the effect of the type of the contours was very similar in both monoscopic and stereoscopic viewing suggests that binocular and monocular mechanisms for 3D shape reconstruction from contours involve similar mechanisms. Specifically, both monocular and binocular processing of 3D shapes involves *a priori* constraints such as symmetry of surface and contours, planarity of contours, as well as geodesic constraint. The operation of these constraints in monocular vision was described by Stevens (1981, 1986) and by Knill (1988), and in binocular vision by Stevens & Brookes (1988) and by Mitchison (1988). More recent studies on the role of constraints in monocular and binocular shape perception include Pizlo, Li & Chan (2005), Pizlo, Li & Francis (2005) and Chan et al. (2006).

Despite similarities of the patterns of results, stereoscopic performance was substantially higher than the monoscopic one. This large difference might, at least in part, be explained by the fact that in stereoscopic viewing, the subjects could see not only the front part of the surface but also the back part.

Consider now the role of geodesic constraint. Although this constraint can account for the differences between vertical and oblique planes, it cannot easily account for the high performance with spheres because the contours produced by spheres were neither planar nor geodesic. There is another explanation for the effect of the types of the contours. In the case of vertical planes, as well as spheres, the symmetry of the contours themselves reflected (at least approximately) the symmetry of the Gaussian surface. This was not the case with oblique surfaces. Contours produced by oblique surfaces suggested that the surface is symmetric along the oblique direction, whereas the surface was actually symmetric along the vertical direction. This explanation based on symmetry can more easily account for the fact that performance with contours produced by spheres was quite high and similar to that with vertical planes. The role of symmetry as a constraint in 3D shape reconstruction was demonstrated by Hochberg & McAlister (1953), Attneave & Frost (1969), and was used explicitly in the computational model of Pizlo and his colleagues (Pizlo, Li & Chan, 2005; Chan et al., 2006). Symmetry seems to be a more robust and reliable constraint than the geodesic constraint. To shed more light on the role of the geodesic constraint vs. that of symmetry requires further studies, both computational and psychophysical.

The Perspecta volumetric display is a technology still at its infancy. While it offers the great advantage of natural stereoscopic viewing, it also has important disadvantages when compared to the more mature LCD technology, such as reduced color resolution, low image brightness, the lack of ability to display opaque surfaces, and wobbling due to the imperfect mechanical rotation of the screen. Therefore, in addition to studying the mechanisms involved in perception of 3D shapes, the work presented here has a second important role, namely the evaluation of volumetric display technology.

We have previously reported [Rosen 2004] the results of another comparison between shape perception on Perspecta and on LCD. The stimuli used then consisted of complex objects with rich texture (e.g. buildings, cars). The results of that earlier experiment did not favor the volumetric display. The different results are explained by the different nature of the stimuli used. The complex shapes and textures exacerbate the limitations of the Perspecta. The drastically simplified 3D image deprives the subject from many cues which collectively outweigh the advantage of stereoscopic viewing. In the present experiment, the stimuli were little affected by the limitations of the volumetric display, and did not erode the advantage of stereoscopic viewing.

6. FUTURE WORK

We will continue to study 3D shape perception along the two intertwined display independent and display dependent directions. The volumetric display is just one of many options available for stereoscopic viewing, and rigorously identifying the strengths and weaknesses of each technology is long overdue. On the other hand, results of fundamental research on 3D shape perception will lead to the advancement of 3D display technology by suggesting perceptually effective resource allocation and approximations.

ACKNOWLEDGEMENTS

The authors thank Yunfeng Li for performing statistical analyses. We would like to thank all the members of our computer graphics and visualization laboratory that put up with the dim lighting conditions required to run the numerous experiments. This research has been supported in part by NSF grants DMS-0138098, EEC-0227828, EIA-0216131, and ACI-0325227. Hoffmann is also supported in part by an IBM Faculty Award.

REFERENCES

Actuality Systems. Perspecta Display. http://www.actualitysystems.com/site/content/ perspecta_display1-9.html.

E. ALLGOWER AND S. GNUTZMANN. "Simplicial pivoting for mesh generation of implicitly defined surfaces.". Computer Aided Geometric Design 8, 305-325, 1991.

ATTNEAVE F. & FROST R. (1969) The determination of perceived tridimensional orientation by minimum criteria. Perception & Psychophysics, 6, 391-396.

CHAN M.W., STEVENSON A.K., LI Y. & PIZLO Z. (2006) Binocular shape constancy from novel views: the role of a priori constraints. Perception & Psychophysics (in press).

HOCHBERG J. & MCALISTER E. (1953) A quantitative approach to figural 'goodness'. Journal of Experimental Psychology, 46, 361-364.

J. BLOOMENTHAL. "An implicit surface polygonizer." Graphics Gems IV, P. HECKBERT, Ed., 324-349, Academic Press, Boston 1994.

KNILL D.C. (1992) Perception of surface contours and surface shape: from computation to psychophysics. Journal of the Optical Society of America, A, 9, 1449-1464.

KNILL D.C. (2001) Contour into texture: information content of surface contours and texture flow. Journal of the Optical Society of America, A, 18, 12-35.

MACMILLAN N.A. & CREELMAN C.D. (2005) Detection theory: a user's guide. Mahwah, NJ: Erlbaum.

MITCHISON G. (1988) Planarity and segmentation in stereoscopic matching. Perception, 17, 753-782.

PIZLO Z. (2001) Perception viewed as an inverse problem. A mini-review. Vision Research, 41, 3145-3161.

PIZLO, Z., LI, Y. & CHAN, M.W. (2005) Regularization model of human binocular vision. Proceedings of IS&T/SPIE Conference on Computational Imaging, vol.5674, pp. 229-240.

PIZLO, Z., LI, Y. & FRANCIS, G. (2005) A new look at binocular stereopsis. Vision Research, 45, 2244-2255.

PAUL ROSEN, CHRISTOPH HOFFMANN, VOICU POPESCU, ZYGMUNT PIZLO. Perception of 3D spatial relations in 3D images. IS&T/SPIE 16th International Symposium Electronic Imaging: Science and Technology 2004 (8 pages).

STEVENS K.A. (1981) The visual interpretation of surface contours. Artificial Intelligence, 17, 47-73.

STEVENS K.A. (1986) Inferring shape from contours across surfaces. In: From Pixels to Predicates, A. PENTLAND (Ed.), (Ablex, Norwood, NJ), pp. 93-110.

STEVENS K.A. & BROOKES A. (1988) Integrating stereopsis with monocular interpretations of planar surfaces. Vision Research, 28, 371-386.