

Analysis of Disparity Distortions in Omnistereoscopic Displays

VINCENT COUTURE, Université de Montréal

MICHAEL S. LANGER, McGill University

SÉBASTIEN ROY, Université de Montréal

An omnistereoscopic image is a pair of panoramic images that enables stereoscopic depth perception all around an observer. An omnistereoscopic projection on a cylindrical display does not require tracking of the observer's viewing direction. However, such a display introduces stereo distortions. In this article, we investigate two projection models for rendering 3D scenes in omnistereoscopic. The first is designed to give zero disparity errors at the center of the visual field. The second is the well-known slit-camera model. For both models, disparity errors are shown to increase gradually in the periphery, as visual stereo acuity decreases. We use available data on human stereoscopic acuity limits to argue that depth distortions caused by these models are so small that they cannot be perceived.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Virtual reality; visible line; surface algorithms*

General Terms: Stereo, Perception

Additional Key Words and Phrases: Panorama, median plane, depth acuity

ACM Reference Format:

Couture, V., Langer, M. S., and Roy, S. 2010. Analysis of disparity distortions in omnistereoscopic displays. *ACM Trans. Appl. Percept.* 7, 4, Article 25 (July 2010), 13 pages. DOI = 10.1145/1823738.1823743 <http://doi.acm.org/10.1145/1823738.1823743>

1. INTRODUCTION

Binocular depth perception requires an observer to establish point correspondences between two images, and to use the disparity differences as a cue to relative depth of visible surfaces. In designing binocular displays such as 3D cinema, it is traditionally assumed that the baseline joining the two eyes is known relative to the screen and, in particular, that the baseline is parallel to the screen. Other methods have relaxed these assumptions, though. For example, in some virtual environments such as CAVEs [Cruz-Neira et al. 1992, 1993], a head tracking system has been used which allows the viewer position and the viewer's orientation to be updated. These environments aim to display exact stereo images to a single observer.

Authors' addresses: V. Couture, DIRO, Pavillon André-Aisenstadt, Université de Montréal, Québec, Canada; email: chapelv@iro.umontreal.ca; M. S. Langer, McGill University, 845 Sherbrooke St., W. Montreal, Québec, Canada; S. Roy, DIRO, Pavillon André-Aisenstadt, Université de Montréal, Québec, Canada.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1544-3558/2010/07-ART25 \$10.00

DOI 10.1145/1823738.1823743 <http://doi.acm.org/10.1145/1823738.1823743>

ACM Transactions on Applied Perception, Vol. 7, No. 4, Article 25, Publication date: July 2010.



Fig. 1. Omnistereoscopic image rendered from a 3D scene model of the Charles Church in Plymouth (UK), courtesy of Karol Kwiatek. The image is encoded in red/cyan anaglyph format.

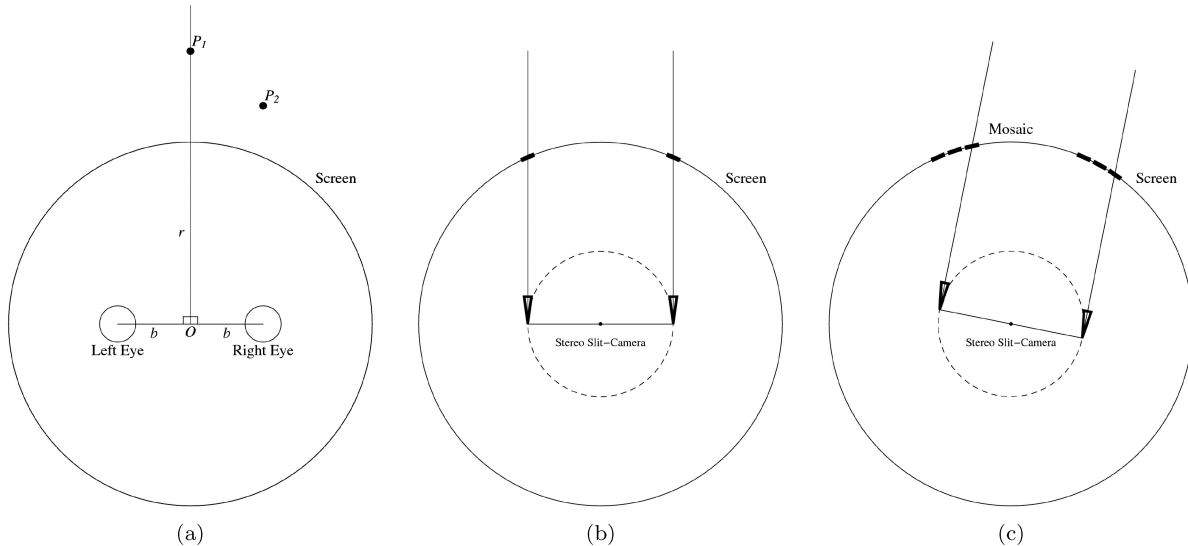


Fig. 2. (a) Omnistereo immersive environment viewed from above: a cylindrical screen of radius r and an observer at the center of the screen. Two scene points are shown. The projection model introduced in Section 3 gives zero stereo disparity error for a point P_1 on the median line. Errors increase in the periphery of the visual field, for example, for a point P_2 . (b) Previous work on omnistereo images uses a rotating stereo pair of slit-cameras verged at a specific distance. Here the distance is infinity, that is, cameras are parallel. (c) Image slits are stitched together in a mosaicing process to cover 360° .

Another approach is to use omnistereoscopic images, which are multiviewpoint panoramic images that contain stereo information all around an observer [Ishiguro et al. 1992; Naemura et al. 1998; Huang and Hung 1998; Peleg et al. 2001; Peleg and Ben-Ezra 1999]. Similarly to CAVEs, omnistereo images can be used for navigation in a virtual environment. However, they remove the need to track the head orientation [Naemura et al. 1998; Bourke 2006, 2009]. An example of an omnistereo image is shown in Figure 1.

Figure 2(a) illustrates an omnistereoscopic display that consists of a cylindrical screen and an observer located at the center O . The baseline of the observer’s eyes is perpendicular to the fixation point, which can be anywhere along the line through O that is perpendicular to baseline, called the “median line”. In this setup, the observer is free to rotate his head, that is, the baseline orientation, but the position of the observer is assumed to remain at or near the center O . We note that the ratio of baseline to display radius is typically much smaller than that illustrated in the figure, so the model is less sensitive to the exact observer position.

One of the challenges of creating omnistereo images is that it is impossible to render correct stereo disparities for all observation orientations at the same time, since the correct rendered stereo disparity depends on the orientation of the observer. In this article, we analyze the distortions that are present in omnistereo displays. We investigate two projection models for rendering omnistereo images from 3D

scenes. For the first model, the disparity errors are *designed* to be zero on the median plane between the eyes regardless of which direction the observer is oriented, and to gradually increase towards the periphery of the visual field [Simon et al. 2004]. This design is motivated by the spatial acuity properties of the human visual system, in particular stereo acuity is highest in the fovea and decreases precipitously with eccentricity. We also investigate the well-known *slit-camera* model and show that it produces similar disparity errors. Moreover, we show that for both models the disparity errors are so small that they are perceptually negligible within a 20° eccentricity. To our knowledge, this is the first attempt to connect depth distortions in omnistereo environments to known stereo vision limits of human observers. Finally, we briefly describe system implementation of the model.

A layout of this article is as follows. In Section 2, we briefly review prior works on omnistereo imaging. In Section 3 we present a projection model that gives zero disparity error for all points on the median plane between the eyes of an observer centered in a cylindrical omnistereo display. The omnistereo distortions caused by this model are discussed in Section 4. Section 5 shows that the standard slit-camera model also causes similar distortions, at least when points are not too far from the vergence point of the cameras. Then, existing limits of human stereo acuity are discussed in Section 6 in which we argue that disparity errors for both models are too small to be perceived. Our discussion is restricted to horizontal disparities only, that is, we do not address vertical disparities. Details on the implementation are presented in Section 7. We conclude in Section 8.

2. PREVIOUS WORK

Most work on omnistereo images addresses how they can be *captured* with a stereo camera [Ishiguro et al. 1992; Naemura et al. 1998; Huang and Hung 1998; Peleg et al. 2001; Peleg and Ben-Ezra 1999]. In Peleg et al. [2001] and Peleg and Ben-Ezra [1999], a stereo pair of cameras is rotated to fully cover 360° degrees (see Figure 2(b) and 2(c)). At every one or two degrees, slit-images are captured having a small horizontal field of view, whose angular width depends on the amount of rotation between consecutive frames. In practice, a set of columns (say 50) is considered for HD images. The resulting omnistereo images are usually displayed as a panorama on a small planar surface such as a monitor [Peleg et al. 2001].

In contrast, this article considers immersive environments in which omnistereo images are rendered and displayed on a cylindrical screen surrounding the viewer. To our knowledge, the only prior published method that uses a cylindrical screen for projecting rendered omnistereo images is Bourke [2006], which uses the preceding slit-camera projection model [Ishiguro et al. 1992; Peleg et al. 2001]. It is observed in Bourke [2006] that the frame of the stereo glasses limits the view window for stereo input to the eyes, but otherwise there is no mention of the disparity information available to the observer and possible perceptual limits. The present article is directly concerned with such perceptual limits. In particular, we investigate the resulting disparity errors to see if the depth distortions they induce are well above known detection thresholds and hence whether they can be perceived.

3. MEDIAN PLANE PROJECTION MODEL

We first describe a projection model that gives zero disparity error for *all* points on the median line between the two eyes. This model is slightly different from the standard slit-camera model in that the latter causes disparity errors on the median plane which depend on the vergence of the eyes (or cameras, in the case of image capture). See Section 5 for details on the slit-camera projection model.

For simplicity, we first present the model for the 2D case (see Figure 3). Recall from Figure 2(a) that P is a point in the scene, and O is the center of the screen circle of radius r . The head is centered at O and each eye is located on a circle of radius b centered at O , such that $2b$ is the *baseline* distance between the eyes. For the cylindrical display in our lab, $r = 230$ cm. For the plots and computations

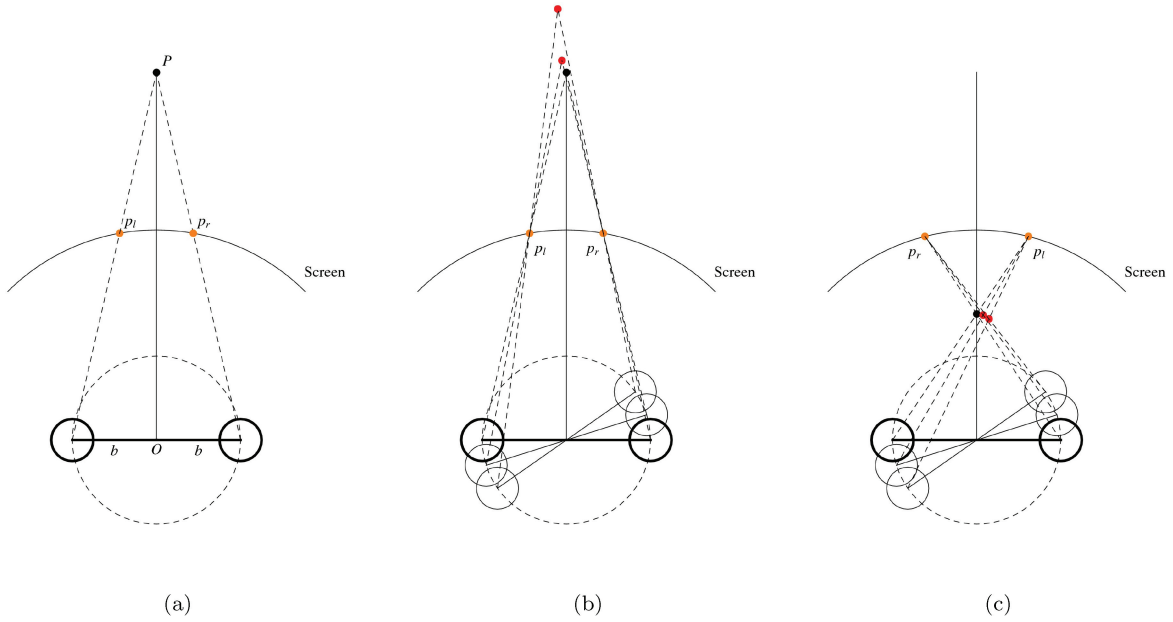


Fig. 3. (a) Rendering positions p_l and p_r of a point P are computed with respect to the baseline orientation for which P lies on the median line between the eyes. Note that only part of the full circular screen is shown. (b) If the observer rotates his head, then P becomes a point in the periphery of the visual field. The rendered points p_l and p_r on the screen are triangulated again and a distortion is introduced. The triangulated P appears at a different location (see red dots). Points behind the screen appear further away in depth. (c) Points in front of the screen appear closer in depth.

later in the article, we take $b = 3.25$ cm. Note that for illustration purposes, Figure 3 uses a larger $b : r$ ratio in than in the actual lab setup.

The projection model requires known scene depths, namely we have a virtual 3D scene model that is being rendered. Given b and r , we compute for each point P the rendered screen positions p_l and p_r , that is, the positions *on the cylindrical display screen* where the rendered point P is projected for the left and right eye's image, respectively. Because it is well known that stereo acuity is highest at the center of the visual field [Howard and Rogers 2002], we design a projection model that gives zero disparity error for a point P when an observer is oriented so that the median line passes through P .

For any point P , we therefore render this point by assuming that the point lies in the head's median plane. Because the display is rotationally symmetric, we consider without loss of generality the eyes located at $(\pm b, 0)$ and a point $P = (0, Z)$. The screen pixel positions p_l and p_r are each computed by intersecting a line with a circle, namely a line joining the corresponding eye and P with the circle of radius r centered at the origin. For the right eye, this intersection is given by

$$(p_{r,x}, p_{r,z}) = \left(\frac{b(Z^2 - \sqrt{\Delta})}{Z^2 + b^2}, \frac{Z(b^2 + \sqrt{\Delta})}{Z^2 + b^2} \right), \quad (1)$$

where

$$\Delta = r^2(Z^2 + b^2) - Z^2b^2. \quad (2)$$

The screen position for the left eye is computed similarly, using $-b$ instead, giving

$$(p_{l,x}, p_{l,z}) = (-p_{r,x}, p_{r,z}). \quad (3)$$

We extend the preceding projection model to the 3D case by considering eyes at $(\pm b, 0, 0)$ and a point $P = (0, Y, Z)$, with the display now a vertical cylinder of radius r centered at the origin. Screen positions $p_{r,x}$, $p_{r,z}$ remain the same, and the vertical screen position is given by

$$p_{l,y} = p_{r,y} = \frac{Y p_{r,z}}{Z}. \quad (4)$$

A point P that lies on the median plane is projected to the correct screen positions p_l and p_r , and so in principle its 3D position can be correctly estimated by triangulation. For 3D points that are not on the median plane, triangulation errors occur that lead to small geometric distortions. The severity of these errors increases gradually with eccentricity. In the following section, we will analyze these errors.

Before doing so, we elaborate on a few assumptions of the projection model. First, when projecting a point P , we are assuming that the observer's eyes are located as in Figure 3 and that the observer is fixating somewhere on the head's median plane. Our analysis does not consider disparity errors relative to stereo acuity when the observer is fixating left or right of the head's median plane. Second, at each new fixation, there is a slight shifting of center of projection (the pupil) as the eyes rotate, since the pupil is slightly displaced from the center of rotation. Since this displacement is so small relative to the baseline, we ignore it in our model. A third assumption is that the model is using a pinhole projection, and so we are ignoring blur and accommodation. As in typical stereo displays, our images are focused on the screen and this leads to a vergence-accommodation conflict [Hoffman et al. 2008]. However, this accommodation conflict is most significant for screens closer than 2 m and so in our setting the conflict would only arise for objects rendered to be closer than the screen.

4. GEOMETRIC DISTORTIONS AND DISPARITY ERRORS FOR MEDIAN PLANE MODEL

In the previous section, we discussed a model for projecting a 3D point onto a cylindrical screen such that a point is triangulated to its correct 3D position when the point lies on the median plane of the observer. For points that are not on the medial plane, triangulation errors occur which lead to small geometric distortions.

Figure 4 illustrates the distortions that are caused by the model when the screen radius is $r = 230$ cm and the eye baseline is $2b = 6.5$ cm. In Figure 4(a) through (c), three planar surfaces are shown in black at different depths, with the distortions shown in red. There is zero distortion at eccentricity 0° , by design. In addition, points on the screen are not distorted at all (see Figure 4(d)). Errors increase gradually away from this zero-distortion locus. At large horizontal eccentricities, depths are distorted to be closer for points in front of the screen, and farther for points behind the screen.

For a 3D point P in the periphery, the rendering positions p_l and p_r generate vertical disparities and do not triangulate to a unique point. Vertical disparities can also arise when the observer is not located at the true center of projection, similar to traditional flat stereo displays [Held and Banks 2008; Woods et al. 1993]. For the plots of Figure 4(a) through (c), the triangulated point was computed using a least squares fit.

In Figure 4(d), black circular curves with varying radius are shown in the horizontal plane passing by the two eyes. The distortions in red are almost entirely in depth, rather than in direction. In particular, black grid lines are drawn every 5° in eccentricity, but they are hidden by the red grid lines because distortion of visual direction is near zero.

Note that points that are far away from the observer undergo larger depth distortions, but these distortions are not necessarily perceivable. The reason is that the visual system measures disparity, which depends on inverse depth. Large absolute errors in triangulated depth might still produce small disparity changes.

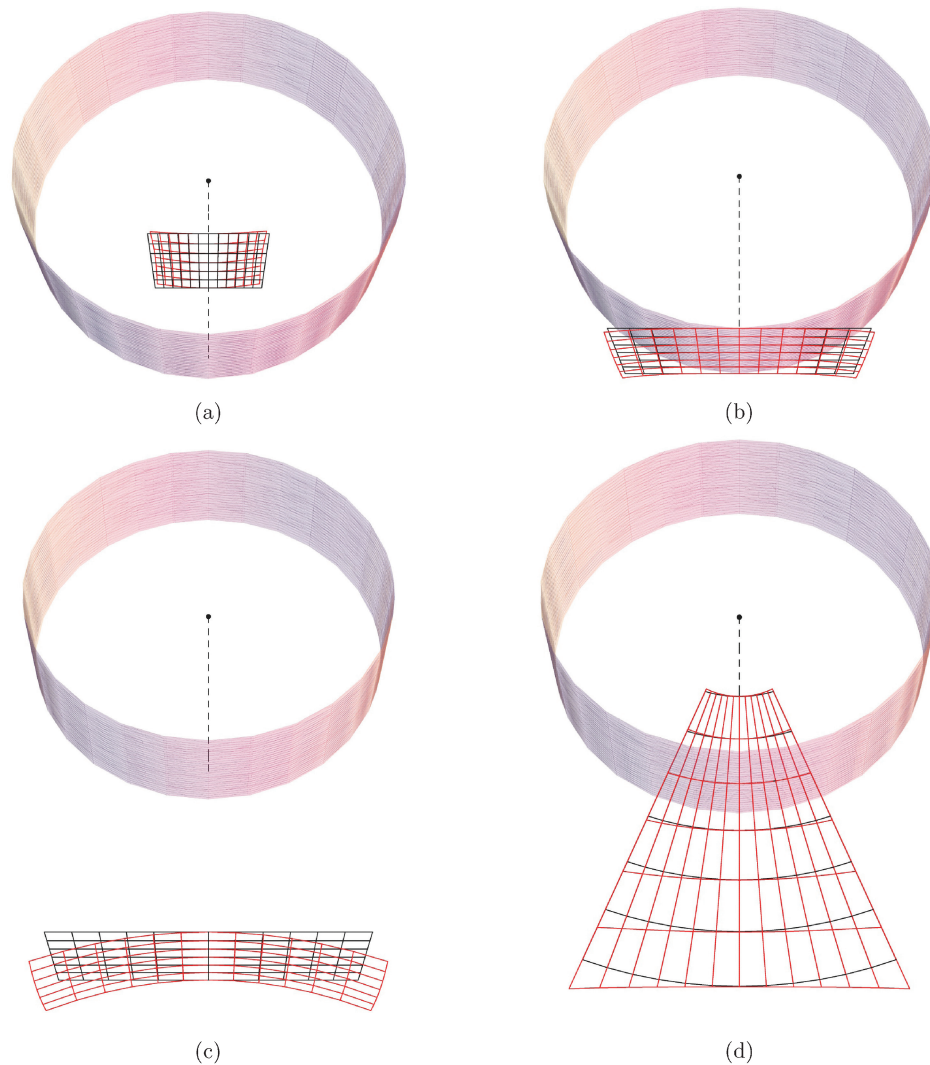


Fig. 4. For an observer looking at a specific orientation on the cylindrical screen indicated by the dashed line, the omnistereo projection model causes distortions in the periphery of the visual field. In (a–c), a plane is shown at different depths, both its true shape (black) and its triangulated distorted shape (red). In (a), the left and right sides of the plane in front of the screen are distorted to be even closer. In (b–c), the sides of the plane are distorted to be farther away. Points on the screen itself are not distorted. Note that these planes are also slightly distorted vertically, but these distortions are very small and will not be discussed further. In (d), circular curves with varying radius (black) are shown in the x - z plane with their perceived distorted shapes (red). Black grid lines that are radial (constant direction) are drawn, but they are not visible because they are overwritten by the red grid lines because distortion of visual orientation is very small.

Figure 5 shows the disparity errors for points within a 40° field of view and at various distances from the observer. The disparity errors are zero at an eccentricity of 0° (all curves) and for points on the screen (i.e. blue curve, 230 cm). The disparity errors are in the range of 0 to 6 arcmin up to 20 degrees in eccentricity, whereas the visual direction distortions are less than 0.1 arcsec (data not shown). Again note that the sign of disparity error is opposite for points in front and behind the screen.

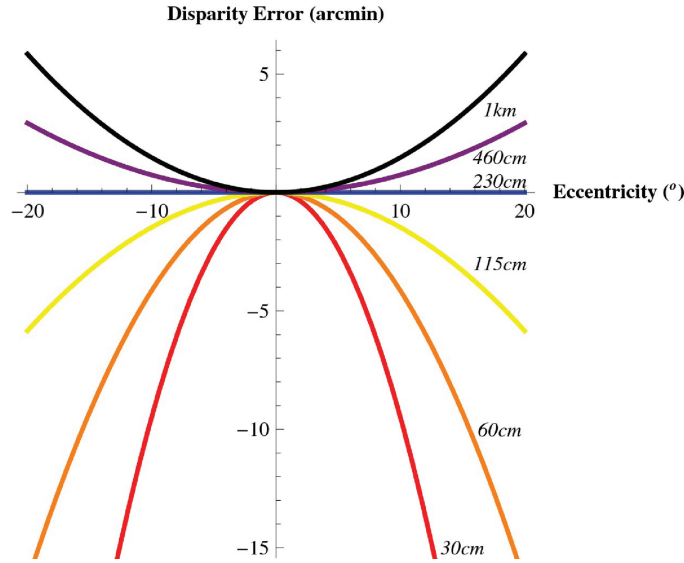


Fig. 5. For an observer looking at a specific orientation, omnistereoscopic projection distorts points in the periphery of the visual field. This plot shows the computed disparity distortions for points within a 40° field of view at various distances from the observer. Distortion is zero at an eccentricity of 0° and for points on the screen.

This is consistent with Figure 3 and plots in Figure 4 which show that points closer than the screen appear even closer and points behind the screen appear farther.

In Section 6, we examine whether these depth errors are perceivable. We consider the disparity errors of the triangulated points and compare these disparity errors to the disparity detection thresholds in human vision.

5. SLIT-CAMERA PROJECTION MODEL

This section compares distortions of the projection model described in Section 3 to the more standard slit-camera model. We assume that the cameras have a baseline of width $2b$ centered at and rotating about O . For simplicity, we also suppose for the remainder of this section that the optical axis of both slit-cameras intersect at infinity, that is, that both cameras are parallel.

Ignoring occlusions, a point P on the median line but not located at infinity is then captured by a V-shaped baseline (see Figure 7(a)). In practice, this means that P is not captured in the same left and right stereo frame. This creates distortion even for points on the median line, in contrast to the median line model presented in Section 3(b) (see Figure 7(b)).

As a projection model, the screen position p_r is computed by intersecting a line joining the right eye $(b \cos(\alpha), b \sin(\alpha))$ and $P = (0, Z)$, and the circle of radius r centered at the origin. Angle α is given by

$$\alpha = \frac{\pi}{2} - \arccos\left(\frac{b}{Z}\right). \quad (5)$$

Figure 6 shows the disparity errors. Observe that the errors are large for very close points to the observer even for an eccentricity of 0° . However, if scene points are limited to points further away than 60 cm, the slit-camera model gives near-zero disparity errors.




Fig. 6. Disparity errors for the slit-camera model for points in a 40° field of view at various distances from the observer. Similarly to the median line model, errors are zero for points on the screen (230 cm). However, very near points may be distorted at some distances even if eccentricity is 0° . (In this plot, the eyes are converging at infinity.)

6. DISPARITY ERRORS VERSUS STEREO ACUITY

On the one hand, since geometric distortions increase with eccentricity, one might expect these distortions to be perceivable at large eccentricities. On the other hand, since the resolution of the visual system decreases with eccentricity, one might expect the distortions not to be perceived. This raises the question of how large the distortions are in comparison to known visual stereo acuity limits, especially in the periphery.

While it is generally agreed that human stereo vision is worse in the periphery, relatively little is known about how performance falls off with eccentricity. Most studies of stereo in peripheral vision only consider eccentricities up to about 10 degrees [Howard and Rogers 2002], and classical experiments consider only very simple local tasks such as depth discrimination of thin isolated vertical lines.

More recent studies have measured perception of more global properties of scene geometry, namely sensitivity to disparity corrugations in the periphery. For example, Prince and Rogers [1998] tested eccentricities up to 20 degrees and, for each eccentricity, they measure the detection threshold of sinusoids of disparity corrugation. They used short presentation times (500 ms), and a fixation point at the center of an annulus that was itself filled with a random dot pattern. They found that peak sensitivity to corrugations was a bandpass function and that for greater eccentricities, peak sensitivity occurred at lower spatial frequencies of the disparity corrugation. The peak detection thresholds themselves increased with eccentricity. For 0, 3.5, 7, 13, and 21 degrees eccentricity, the peak thresholds were about 0.03, 0.3, 0.5, 2, and 5 arc minutes of disparity, respectively. We emphasize that “peak” here refers to the corrugation spatial frequency that was most easily detected.

The peak thresholds just mentioned are quantitatively similar to the distortions in Figure 5 and Figure 6, provided that the rendered points are at least 1m away from the viewer (yellow curve and above). Of course, one should not attach too great a significance to the similarities of the data; the psychophysical data are dependent on the details of the experiment (observers’ task, stimulus, presentation time, definition of threshold, etc.). Nonetheless, the similarities do suggest that the geometric

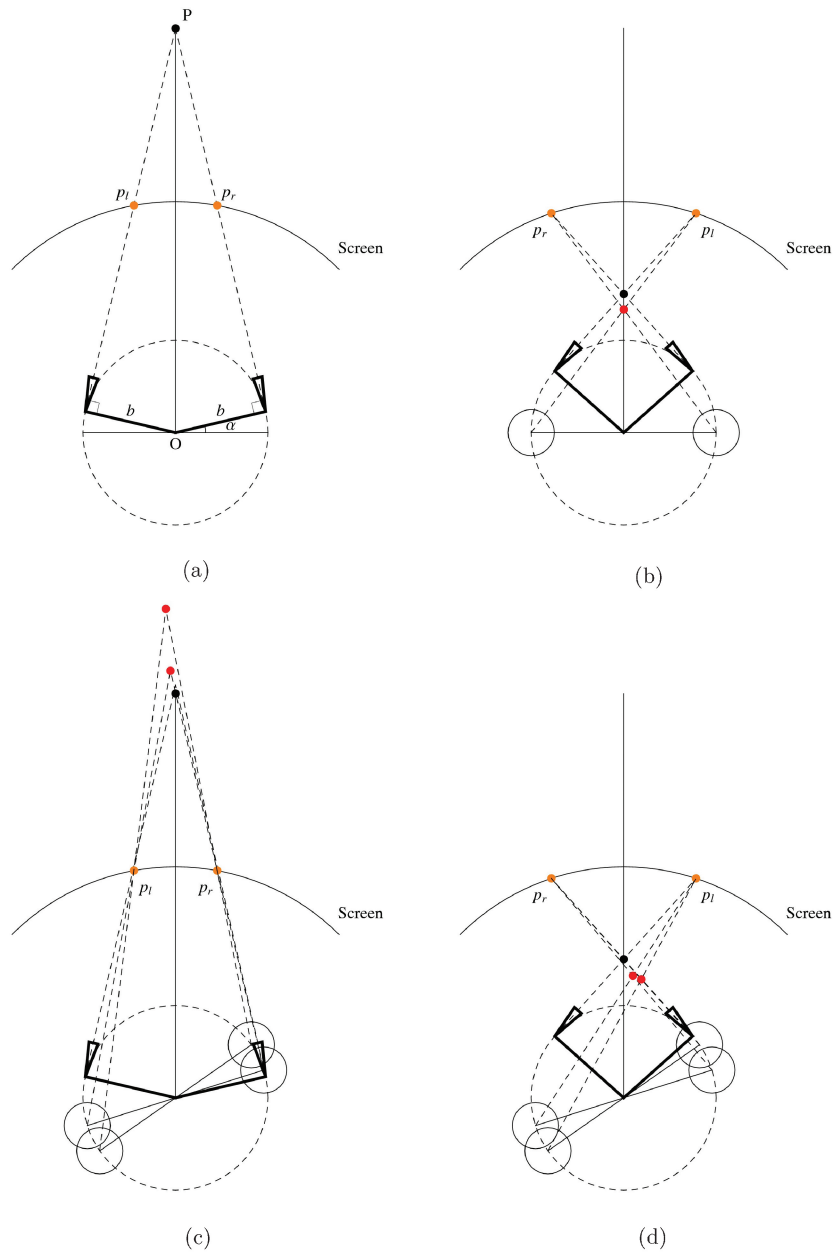


Fig. 7. (a) The slit-camera projection model projects a point P on the circular screen with respect to the position of the left and right cameras having optical axis intersecting P . (b) Point P may appear distorted even if it lies on the median line of the observer. (c–d) Similarly to Figure 3(b–c), P appears at a different location (see red dots) if in the periphery.

distortions introduced by both projection models are at or below the detection threshold, and hence may not be of significance to human observers.

As an aside, the reader also may be interested in knowing why human stereo acuity worsens in the periphery. There is evidence that the high detection thresholds mentioned earlier are due mainly to

the lower resolution of the luminance signal in the left and right eye image, rather than to limitation in stereo processing per se. The luminance signal decreases in the periphery because of factors such as poorer optics, sparser retinal sampling, and greater pooling of photoreceptors by each ganglion cell [Banks et al. 1991]. Stereo performance worsens in the periphery but no worse than one would expect from the worsening input luminance signals [Hess et al. 1999; Banks et al. 2004]. Indeed, at sufficiently low luminance spatial frequencies (or, equivalently, low dot densities if one is using random dot stereograms) detection thresholds for disparity corrugations in the fovea are similar to those in the periphery. See Figure 6 of Banks et al. [2004], for example.

Finally, we should mention that our discussion of stereo acuity in human vision is far from complete. In particular, we have considered acuity limits on horizontal disparities only. Vertical disparities are also often present and appear to be treated differently than horizontal disparities by the human visual system, for example, in the pooling of information across the visual field [Adams et al. 1996]. A more complete study of the effects of disparity distortions should consider vertical disparities. For recent reviews of some of the relevant literature, see, for example Held and Banks [2008] and Read and Cumming [2010].

7. IMPLEMENTATION

This section describes a rotation method that leads a point P to be rendered at the screen positions p_l, p_r . The method is similar in flavor to Simon et al. [2004] and can be applied with both the median plane and the slit-camera projection models.

We first present the method in 2D. As shown in Figure 8 for the right eye, the method rotates P by θ to get P_r , which is rendered with respect to the origin O . Note that angle θ is of opposite sign for the left eye. Screen points p_l and p_r coincide ($\theta = 0$) for scene points P that lie on the screen. Angle θ can be computed by

$$\theta = \arctan(p_{r,z}, p_{r,x}). \quad (6)$$

The 3D extension is a rotation θ within the epipolar plane defined by the two eyes and point P .

We implement the method using a vertex shader that rotates each vertex by its corresponding θ , using a positive rotation for the left view and a negative rotation for the right view. For the setup in our lab and for points farther than 100 cm from the observer, the magnitude of rotation θ is less than 1° . Note that this is much less than the rotational shears that are shown in Figure 8, where the baseline is exaggerated.

The two sheared scenes, that is, for the left and right eyes, are each rendered such that the center of projection is at O . In our vertex shader implementation, the diffuse reflection term is computed using the unmodified vertex and light positions and normals. The specular term for point P is computed by assuming the head is oriented such that P is on the medial plane.

Note that because a vertex shader is applied on vertices and not pixels, the rotational model tends to distort long edges in low tessellated scenes, as only the endpoints (the vertices) are moved correctly. Hence, the rotational model works best for a highly tessellated scene.

One caveat is that we are assuming that any point P that is visible to the both eyes for the real observer will also be visible in both of the rendered images. Our implementation does not guarantee this condition is met, however, since a point at another depth and off the optical axis could in principle be rotated such that it occludes point P in one of the two images.

Finally, we note that the projection model of Section 3 can be implemented in other ways. For example, rather than rotating about O , a translation parallel to the baseline could be used. This would

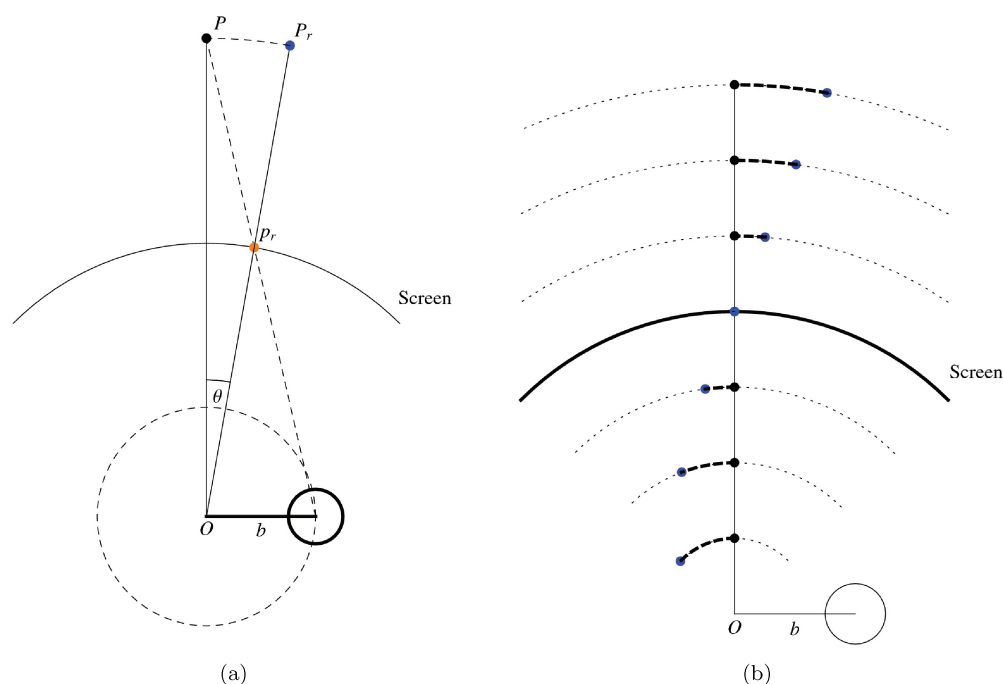


Fig. 8. (a) Left and right views are rendered from O . Disparities between the two views are created by modifying the position of scene points by a rotation around O . The magnitude of rotation depends on the distance from the eye to the point. Sign of rotation is opposite for the two eyes. The angle of rotation θ is given in Eq. (6). Baseline b is typically much smaller than screen diameter and as a result rotations are typically less than 1° . (b) Each point P is relocated by a rotation on a circle of radius $|P|$, and rendered with respect to O . These rotations roughly shear the scene in cylindrical coordinates. This shear produces zero distortion if P is on the observer's medial plane.

give rise to a shearing of medial plane that is parallel to the baseline. Since rotations are typically small (less than 1°), this new shear would be near identical to the one produced by rotation. Hence the distortions would be similar as well.

The rotational projection model was tested in a 230 cm radius cylindrical screen, with a height of 150 cm. Four projectors were used to cover half the screen (180 degrees), with neighboring projectors overlapping. Lighttwist [Tardif and Roy 2005; Tardif et al. 2003], an open-source multiprojector system, automatically aligns the projectors from the point of view of a camera, here at the center of the cylinder screen, without actually reconstructing the screen in 3D as in Raskar et al. [1999]. High pixel resolution and contrast we achieved at an affordable cost by the use of HD projectors.

For polarized stereo projection, the number of projectors is doubled to eight. The light of the projectors for the left eye is polarized horizontally, and vertically for the right eye. A special screen maintains light polarization, and observers must wear appropriate filtering glasses. Real-time navigation was also successfully achieved by having a rendering computer connected to each projector, synchronized by a master computer that multicasts the joystick input. Navigation was controlled by a single observer.

In practice, the observer might be located off-center, especially if more than one observer is allowed in the omnistereo environment. In this case, perspective distortions arise when the viewer is far from the assumed center of projection [Vishwanath et al. 2005; Banks et al. 2009].

8. CONCLUSION

This article presented a projection model for rendering omnistereo images from 3D scenes such that the 3D distortions of the scene are zero for points P that are in the center of the field of view, that is, on the head's median plane. The method assumes the observer is standing at the center of the cylindrical screen. 3D distortions were computed and compared to available stereo acuity measurements. The disparity errors from the projection model were found to be near threshold for detection of disparity corrugation at all eccentricities up to 20 degrees. Future work will analyze perspective distortions that result in omnistereo 3D cinema when the viewer is far from the assumed center of projection.

REFERENCES

- ADAMS, W., FRISBY, J., BUCKLEY, D., GARDING, J., HIPPLISLEY-COX, S., AND PORRILL, J. 1996. Pooling of vertical disparities by the human visual system. *Percept.* 25, 2, 165–176.
- BANKS, M., HELD, R., AND GIRSHICK, A. 2009. Perception of 3-d layout in stereo displays. *Inform. Display* 25, 1, 12–16.
- BANKS, M. S., GEPSHTEIN, S., AND LANDY, M. S. 2004. Why is spatial stereoresolution so low? *J. Neurosci.* 24, 9, 2077–2089.
- BANKS, M. S., SEKULER, A. B., AND ANDERSON, S. J. 1991. Peripheral spatial vision: Limits imposed by optics, photoreceptors, and receptor pooling. *J. Optical Soc. Amer.* 8, 11, 1775–1787.
- BOURKE, P. 2006. Synthetic stereoscopic panoramic images. In *Proceedings of the International Conference on Virtual Systems and MultiMedia (VSMM)*. Lecture Notes in Computer Science, vol. 4270, 147–155.
- BOURKE, P. 2009. Omni-Directional stereoscopic fisheye images for immersive hemispherical dome environments. *Comput. Games Allied Technol.* 136–143.
- CRUZ-NEIRA, C., SANDIN, D. J., AND DEFANTI, T. A. 1993. Surround-Screen projection-based virtual reality: The design and implementation of the cave. In *Proceedings of the ACM Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. 135–142.
- CRUZ-NEIRA, C., SANDIN, D. J., DEFANTI, T. A., KENYON, R. V., AND HART, J. C. 1992. The cave: Audio visual experience automatic virtual environment. *Comm. ACM* 35, 6, 64–72.
- HELD, R. T. AND BANKS, M. S. 2008. Misperceptions in stereoscopic displays: A vision science perspective. In *Proceedings of the 5th Symposium on Applied Perception in Graphics and Visualization (APGV)*. ACM, New York, 23–32.
- HESS, R. F., KINGDOM, F. A. A., AND ZIEGLER, L. R. 1999. On the relationship between the spatial channels for luminance and disparity processing. *Vis. Res.* 39, 3, 559–568.
- HOFFMAN, D., GIRSHICK, A., AKELEY, K., AND BANKS, M. 2008. Vergence-Accommodation conflicts hinder visual performance and cause visual fatigue. *J. Vis.* 8, 3, 33.
- HOWARD, I. P. AND ROGERS, B. J. 2002. *Seeing in Depth*. Oxford University Press.
- HUANG, H.-C. AND HUNG, Y.-P. 1998. Panoramic stereo imaging system with automatic disparity warping and seaming. *Graph. Models Image Process.* 60, 3, 196–208.
- ISHIGURO, H., YAMAMOTO, M., AND TSUJI, S. 1992. Omni-Directional stereo. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2, 257–262.
- NAEMURA, T., KANEKO, M., AND HARASHIMA, H. 1998. Multi-User immersive stereo. In *Proceedings of the IEEE International Conference on Image Processing 1*, 903.
- PELEG, S. AND BEN-EZRA, M. 1999. Stereo panorama with a single camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 1*, 1395.
- PELEG, S., BEN-EZRA, M., AND PRITCH, Y. 2001. Omnistereo: Panoramic stereo imaging. *IEEE Trans. Pattern Anal. Mach. Intell.* 23, 3, 279–290.
- PRINCE, S. J. D. AND ROGERS, B. J. 1998. Sensitivity to disparity corrugations in peripheral vision. *Vis. Res.* 38, 17, 2533–2537.
- RASKAR, R., BROWN, M. S., YANG, R., CHEN, W.-C., WELCH, G., TOWLES, H., SEALES, B., AND FUCHS, H. 1999. Multi-Projector displays using camera-based registration. In *Proceedings of the IEEE Visualization Conference*.
- READ, J. AND CUMMING, B. 2010. Does depth perception require vertical-disparity detectors? *J. Vis.* 10, 6.
- SIMON, A., SMITH, R. C., AND PAWLICKI, R. R. 2004. Omnistereo for panoramic virtual environment display systems. In *Proceedings of the IEEE Virtual Reality Conference*. 67.
- TARDIF, J.-P. AND ROY, S. 2005. A mrf formulation for coded structured light. In *Proceedings of the International Conference on 3-D Digital Imaging and Modeling*. 22–29.
- TARDIF, J.-P., ROY, S., AND TRUDEAU, M. 2003. Multi-Projectors for arbitrary surfaces without explicit calibration nor reconstruction. In *Proceedings of the International Conference on 3-D Digital Imaging and Modeling*. 217–224.

- VISHWANATH, D., GIRSHICK, A., AND BANKS, M. 2005. Why pictures look right when viewed from the wrong place. *Nature Neurosci.* 8, 10, 1401–1410.
- WOODS, A., DOCHERTY, T., AND KOCH, R. 1993. Image distortions in stereoscopic video systems. *Proc. SPIE: Stereoscopic Display Appl. IV*. Vol. 1915. 36–48.

Received June 2010; revised June 2010; accepted June 2010