Draft manuscripts are due to your session coordinator no later than **11 June, 2010**.

Your final manuscript must be sent to Joel Welch at jwelch@smpte.org by **25 June, 2010** to be included in the SMPTE Online Proceedings.

Author(s)

First Name	Middle Name	Surname	Role	Email	SMPTE Member ?
Laurie	M.	Wilcox	1st author	lwilcox@yorku.ca	yes

Affiliation

Organization	Address	Country
York University, Centre for Vision Research, Department of Psychology	CSEB 0003H, York University, 4700 Keele st., Toronto, Ontario	Canada

Author(s) – repeat Author and Affiliation boxes as needed—

First Name	Middle Name	Surname	Role	Email	SMPTE Member ?
Inna		Tsirlin	Co-author	itsirlin@yorku.ca	no

Affiliation

Organization	Address	Country
York University, Centre for Vision Research, Department of Psychology	CSEB 0001, York University, 4700 Keele st., Toronto, Ontario	Canada

Author(s) – repeat Author and Affiliation boxes as needed—

First Name	Middle Name	Surname	Role	Email	SMPTE Member ?
Robert	S	Allison	Co-author	Allison@cse.yorku.ca	yes

Affiliation

Organization	Address	Country
York University, Centre for Vision Research, Computer Science & Engineering	CSEB, York University, 4700 Keele st., Toronto, Ontario	Canada

Publication Information

Pub ID	Pub Date
2010 SMPTE 3D Conference Paper	

SMPTE Meeting Presentation

Sensitivity to monocular occlusions in stereoscopic imagery: Implications for S3D content creation, distribution and exhibition.

Laurie M. Wilcox, Ph.D.

York University, Centre for Vision Research, Department of Psychology, York University, 4700 Keele St., Toronto, Ontario, M3J 1P3, lwilcox@yorku.ca

Inna Tsirlin, M.A.

York University, Centre for Vision Research, Department of Psychology, York University, 4700 Keele St., Toronto, Ontario, M3J 1P3, itsirlin@yorku.ca

Robert S. Allison, Ph.D.,

York University, Centre for Vision Research, Department of Computer Science and Engineering, York University, 4700 Keele St., Toronto, Ontario, M3J 1P3, allison@cse.yorku.ca

Written for presentation at the

SMPTE International Conference on Stereoscopic 3D for Media and Entertainment

Abstract. *Since S3D requires two views of a scene, one for each eye, transformations such as rescaling, 2D to S3D conversion, synthesis of multiview displays, coding and ADAT communications efficiency require generation of new views from 2D images. One of the main challenges to this process is the identification and treatment of monocularly occluded regions. In natural environments, monocular occlusions occur whenever objects are partially obstructed by other objects in a scene, giving rise to a region that is visible to only one eye. Experiments have shown that these regions influence depth percepts. Importantly, if monocular occlusion regions are presented with texture that is inconsistent with the surrounding regions, or with inappropriate geometry, depth is degraded. This paper will review the geometric basis of monocular occlusions and their role in natural depth perception. The analysis will be framed in the context of the reconstruction of novel and appropriate viewpoints from sequences of 2D images from one or more vantage points.*

Keywords. S3D, monocular occlusions, 2D to 3D conversion, stereoscopic depth perception

Introduction

Stereoscopic depth from binocular disparity

The basis for stereoscopic depth perception and thus stereoscopic three-dimensional displays (S3D) is the fact that we have two forward facing eyes that are laterally displaced in the head. At any point in time, the brain receives two similar, but not identical, images of the visual scene. In one type of interocular difference, a feature in the scene may project images on the two retinas that differ in their position relative to the fovea. As shown in Figure 1, this positional disparity, or binocular disparity is registered by the brain, and used to signal the depth of an object relative to where the observer is fixating. Note that in computing binocular disparity, the brain must be able to determine which features or regions of the images correspond in the two eyes. Research in human stereopsis has shown that the brain performs this feat remarkably well, and with extreme precision (for a detailed review see [1])

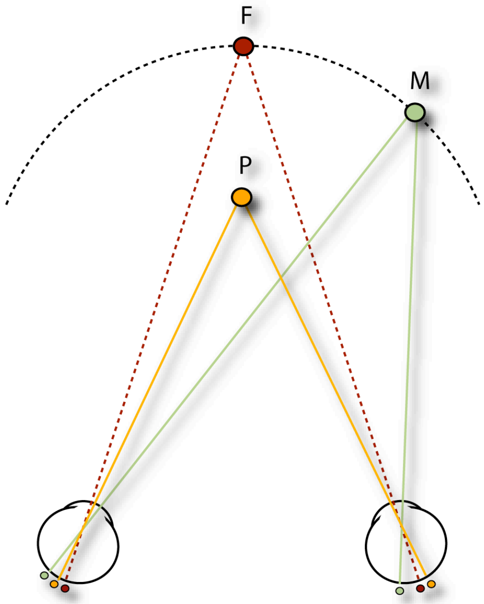

Figure 1. Here we illustrate the geometry of stereopsis as viewed from above. The observer fixates point F, the image of which will fall on the fovea of each eye (indicated in red). Note that point M is located at the same distance from the observer as fixation (more precisely on the Veith-Muller circle). Images of point M will fall in the same direction (here to the left) and distance from the fovea in each eye (shown in green). Point P however is closer to the observer than fixation. Images of P (shown in yellow) will fall on opposite sides of the fovea, and this difference in position is the positional disparity encoded by the stereoscopic system. The coloured dots at the back of the eyes show the relative positioning of the images of F, P and M on the retinae.

Monocular Occlusions

Naturally occurring monocular occlusions

Stereopsis from positional disparity has been the most studied difference between the two eyes' images but other differences are also important. Due to the lateral separation of our eyes and the occlusion of surfaces by nearer objects there are areas in a scene that are visible to one eye only. This fact was noted by Leonardo da Vinci, who observed that because each eye sees slightly more of one side of a solid object, no 2D representation can fully recreate a 3D scene, unless that 3D scene was viewed with one eye [2]. The two most common types of monocular occlusion in the natural environment occur when (1) an object occludes part of a background surface or another, more distant, object (background occlusion) and (2) a volumetric object occludes part of itself (self-occlusion). As shown in Figure 2, the size of occluded background regions depends critically on the separation between the eyes (interocular distance), and the depth between the occluding object and the background. The same dependencies are also true of self-occlusions, where the occluded region is part of a volumetric object. For a more complete review of monocular occlusions see the recent review by Harris and Wilcox [3].

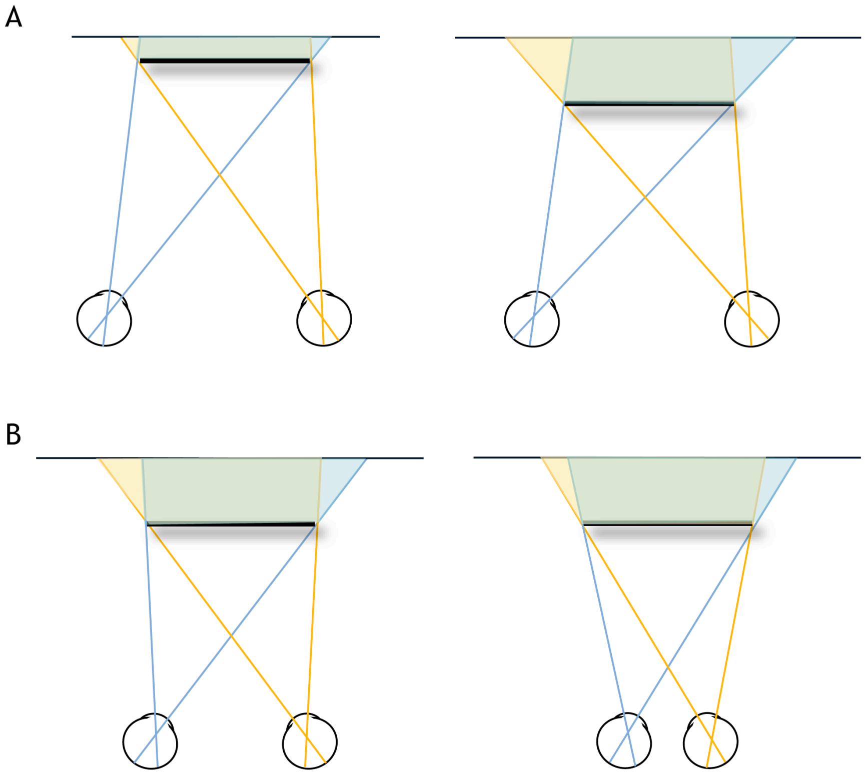

Figure 2. These diagrams depict a plan view of an opaque occluding surface positioned in front of a background plane. The region shown in green represents the background area that is not visible to either eye. Monocularly occluded areas on either side of the central region are hidden from either the left (blue) or the right (yellow) eye. In the top row (A) the occluder is closer to the observer in the rightmost image and the monocular regions are larger. As illustrated in the bottom row (B) the size of the occluded regions also varies with interocular separation, for as the eyes move closer together, the monocular regions become smaller and vice versa. Note that there is also a corresponding increase in the area of the background hidden from both eyes as the interocular distance decreases.

Monocular features arise in binocular vision for reasons other than the two primary sources listed above. For instance there are inherent geometric and physiological differences in monocular visual fields. Some parts of the visual field are seen monocularly at all times, for instance the area corresponding to the blind spot at the optic disc of the fellow eye or the region of space occluded by the nose or brow in one eye but not the other. These are special cases of monocular occlusion by one's own body and human vision has adapted to handle these occlusions without perceptual disruption. Specularities are another class of monocular feature; while binocularly viewed matte objects project similar images in the two eyes, shiny objects can project images with very different intensity or colour in the two eyes. Such differences contribute to the binocular perception of surface gloss. However, interocular mismatches due to reflections can also be troublesome, particularly when they cannot be associated with a surface. Monocular camouflage is a special case of monocular occlusion where a foreground object occludes a similarly coloured or textured background. This makes the occluding, rather than occluded, object effectively invisible in one eye. Finally, a large binocular disparity can result in loss of sensory fusion (diplopia) and eventual loss of depth percepts. Under diplopic conditions, each eye's 'half-image' of the object may be treated by the visual system as a monocular element. In this situation, and others where a feature in one eye is dissociated from its match in the fellow eye, the feature could potentially give rise to false monocular occlusion percepts.

Apart from physiological sources of occlusion, the perceptual issues surrounding these less common types of monocular feature are similar to those resulting from background occlusion due to interposition of a near object or self-occlusion. Therefore we will focus our discussion on these latter sources.

Monocular occlusions in S3D content

In the natural environment monocular occlusions are ubiquitous, as they will occur at the left and right edges of any object at a near viewing distance. When S3D content is captured stereoscopically with two cameras placed at appropriate locations the monocularly occluded regions described above will be registered appropriately, and the resultant S3D content will be fully consistent with the binocular disparity in the scene. However, there are a number of situations in which monocular regions become problematic either because they (i) are inadvertently added, (ii) are distorted, or (iii) must be simulated.

1. Accidental monocular regions.

Examples of the first category include noise from external sources (e.g. dirt, scratches, lens flare, transmission errors) that contaminate the footage for one eye, but not the other. Digital cinema has improved the reliability of stereoscopic production and exhibition, reducing many of the historically troublesome artifacts inherent in analogue medium such as film. Even if the original content is free of artifact, lossy compression techniques for storage and transmission can disrupt or introduce monocular features. Similarly compression algorithms can cause pixels in one eye's view to differ in luminance from the other eye.

2. Distortion of monocular regions

The second category includes situations where appropriate monocular occlusions are present but distorted, for instance they are at the wrong scale or location in the scene. Perceived depth can be distorted under such conditions, a phenomenon which may be particularly problematic in scenes containing self-occlusions. In day-to-day viewing self-occlusion for volumetric objects is most pronounced with small objects directly in front of the viewer. As image features move eccentrically away from straight ahead, the geometric range of self-occlusion (slant of the occluded surface) becomes increasingly constrained. For instance, the left face of a cube, located squarely and directly in front of the observer with width less than the interocular separation will be seen only by the left eye; conversely the right face will be seen only by the right. Imagine a stereo sequence of such a cube blown up to a 30' screen so that the frontal face of the cube is now larger than the interocular distance. Even if convergence and disparity are constrained for comfortable stereopsis the monocular occlusions will no longer be consistent with viewing a cube (as neither eye should see the sides of a cube wider than the interocular distance). Little is known about how the visual system responds to such situations in complex real-world images but work with simpler stimuli predicts that depth distortion and binocular rivalry may be possible (see [4]). Similar issues can arise at the boundaries of 3D shapes when image transformation techniques are used to warp 2D sequences onto 3D models for dimensionalization. Generally any process that transforms disparities in the S3D content will also transform the size and/or ecological validity of monocular zones in the scene, potentially distorting depth percepts.

3. Simulated monocular regions

In the third category, monocular occlusions are indicated by the geometry of the S3D scene, but must be simulated. In general this occurs whenever a new view must be synthesized. Typical situations where this takes place include S3D capture and coding of S3D content with image-plus-depth techniques (i.e. with lidar), interpolating new views to change stereo baseline in post-production, and interpolation of additional or different views for autostereoscopic displays. The most common case today involves creation of S3D content from 2D footage but the issues are similar in any case where a novel view must be created. The problem of generating the 'new' eye's view from 2D content arises from the fact that additional cameras will see different portions of objects and backgrounds, and these occlusions and 'disocclusions' will have to be rendered faithfully. The most widely used 2D to S3D conversion approach involves the use of rotoscoping and/or computational algorithms to identify and segment objects from the background. The objects to be positioned in depth are shifted laterally to introduce binocular disparity. When the two video streams are viewed simultaneously (stereoscopically) the object is seen to lie at a different depth plane from the background; the larger the shift the larger the perceived depth of that region.

When a segmented object from one scene is displaced to introduce binocular disparity, a region the size of the lateral offset will be revealed. Note that this blank area will be visible to one eye and not the other, thus it is a background occlusion. Conversely a complementary portion of the background on the other side will be covered by the shifted object creating a monocular occlusion in the other eye. The 'hole' left by the shifted object is the more problematic occlusion as suitable imagery must be created to fill the gap. The task of the content creator is to fill this region in a manner that is consistent with the remainder of the scene. Typically this requires matching the background texture or content that was previously (or will be) visible in this area. If the background is simple, say a blue sky, with little texture variation, this task is straightforward and simple matter of colour matching. Most often however this monocular region will need to be filled with texture appropriate for a complex scene.

Filling monocularly occluded areas correctly is especially important for dynamic sequences. The content of monocular zones in static scenes is constrained mostly by the surrounding texture and to some extent by the viewing geometry (see Figures 2 and 4). As long as the content of the monocular zones does not violate these constraints and is consistent with the scene it is likely to be accepted by the viewer. In dynamic S3D scenes motion parallax cues are present and the content of the occluded zone may be revealed in other frames making it more likely that discrepancies in depth and content may become noticeable. Furthermore, if different texture samples are used to fill the same monocular areas across multiple frames, the observer will experience motion of the texture elements within the monocular region.

Depth from monocular occlusions: implications for S3D content

In the decades following Wheatstone's [5] seminal paper on binocular depth perception, researchers focussed on the information in each eye that was matched to extract the binocular disparity signal. The presence of unmatched regions was considered either inconsequential, or noise in the binocular matching system. This attitude was implicit in most behavioural studies of stereopsis at the time, but was made explicit in the late 1970's with the computational approaches to modeling stereoscopic depth perception [6, 7] . A typical approach was to identify such monocular regions, and then exclude them from further analysis.

Early evidence for the importance of monocular occlusions came from Adolf von Szily at the beginning of the 20th century. He demonstrated that monocular features in silhouette stereograms can induce depth percepts when fused [8] (Figure 3).

Figure 3. Figures in which object shape is defined by monocular occlusions. When the elements in A are fused by crossing the eyes the depth percept is that of a strip of paper with the ends folded towards the observer as illustrated below. When the image are swapped, as in B, and fused in the same manner the strip of paper appears rotated away from the observer as shown below. Notice how the monocular regions on either side of the figure give rise to a percept of a solid surface (modified from figure 4 in [8]).

However, it was not until about 30 years ago that monocular occlusions received empirical attention. In 1967, Lawson and Gulick [9] confirmed von Szily's findings that the presence of monocular regions in a simple figure promoted a percept of depth. Twenty years following Lawson and Gulick's observations, Gillam and Borsting [10] were the first to show empirically that monocular occlusions accelerate stereoscopic depth perception. Numerous subsequent studies with a range of stimuli have confirmed that monocular occlusions alone can elicit percepts of depth and uncovered other important aspects of depth perception from monocular occlusions.

The effect of monocular occlusion properties on depth percepts

The visual system relies on viewing geometry to interpret depth order from disparity and monocular occlusions. Consequently, correct positioning of occluded regions is crucial for veridical interpretation of depth in a scene. In their seminal paper Nakayama and Shimojo [11] used a simple stimulus depicting a rectangle occluding a bar, to show that when the location of the bar was consistent with an occlusion interpretation (i.e. the bar is visible in the right eye to the right of the rectangle) the bar was clearly seen behind the rectangle. However, when the occluded bar was presented to the other eye, depth percepts were lost. Shimojo and Nakayama [12] tested this observation and found that observers' depth percepts in the invalid condition were unstable and rivalrous.

The content of monocular areas is also important for correct interpretation of depth in a scene. Grove and Ono [13] and Grove, Gillam and Ono [14] found that significantly more time was required to perceive stereograms with occluded zones textured differently from the background. In their experiments, when the occluded region could not be interpreted as part of the background or of the occluding surface, perceived depth was degraded. Others have also shown that when the content of the monocular region is consistent with occlusion geometry rivalry is suppressed [15, 4].

In fact, quite large regions of monocular texture can be introduced to a stereoscopic image without loss of stability, if the context is consistent with the presence of an occluder. For example, Forte, Pierce and Lenny [16] showed that stable scene percepts are obtained from images created to mimic looking through a picket fence, with the fence posts perfectly spaced so that each eye sees a different portion of the background. In their experiments the foreground fence was binocular and fused, and the background, even though it could not be matched in the two eyes at any point, was perceived as stable and continuous. They also found that as the orientation of the background texture changed from horizontal to vertical such that there was no continuity from one monocular patch to another, it became harder for the observers to perceive the background lying at a coherent location in depth.

Essentially the same phenomenon was reported by Howard [17], who showed that a stimulus consisting of a set of black disks in one eye, and white disks in the other, when fused stereoscopically, created a percept of a surface with cut-out holes. The percept is that of a surface with apertures through which the observers sees a background surface which is displaced in depth; the phenomenon was named ‘the sieve effect’. It appears that this percept is a more extreme version of the stimulus later used by Forte et al. [16] . In this case the disks have opposite polarity, and prone to produce rivalry and the percept of lustre. Even so, the visual system is able to form a coherent percept because there is a viable occlusion-based explanation.

Relevance to S3D media

The implications of these findings for generating additional stereoscopic views for S3D film or broadcast are clear. If monocular information is added to S3D content, it must be added in a manner consistent with the scene geometry and content. As illustrated in Figure 4, the location of the monocular occlusion depends on the direction of the displacement of the segmented object in depth relative to its original position in the scene. Let’s assume that the original content is considered the left eye’s view, and the altered content is the right eye view (Figure 4, top row). To displace a portion of the scene (an object) towards the audience in depth, then it will have to be moved to the left in the right eye footage. This results in a monocular region along the right side of the shifted object, a region that would normally be occluded from the left eye’s view by the object (Figure 4, middle row).

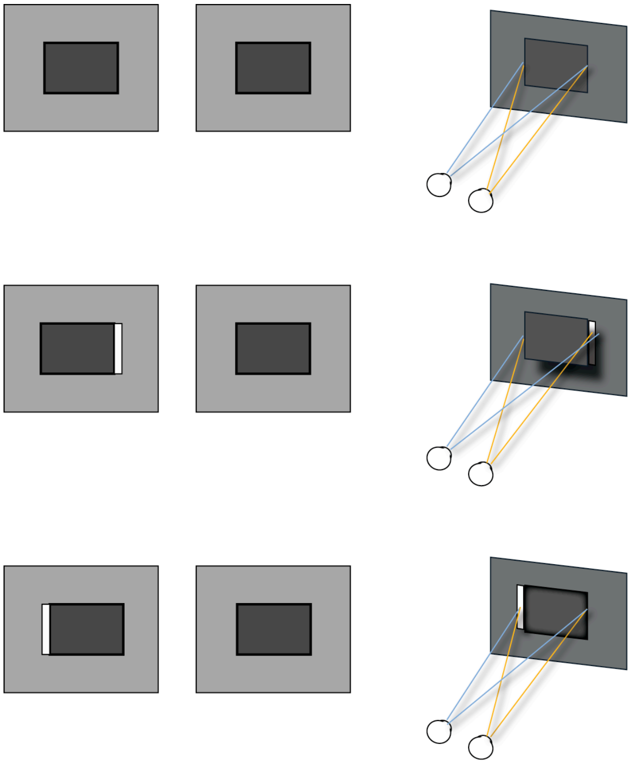

Figure 4. Cross fusion of the first two columns of figures illustrates the depth percepts obtained from each type of disparity and occlusion configuration (note the left eye's view is on the right, and the right eye's view is on the left). The right column shows cartoon images of the percepts in each row. In the top row, the central dark grey rectangle is positioned at the same depth plane as the surround, and no relative depth is perceived. The middle row illustrates a leftward shift of the central region, and the resulting blank area along its right edge. When fused, the dark rectangle is seen with negative parallax, therefore positioned in front of the surrounding rectangle. Even when left white, the percept of the monocular region is consistent with the occlusion relationship (that is, behind the central rectangle). In the bottom row, the central region in the left image is shifted to the right to create a positive parallax. In this case the blank region is 'disoccluded' and though it is monocular, given the viewing geometry it *should* be visible to both eyes (as shown in the diagram to the right). This area will appear unstable unless its texture and luminance is carefully matched to the background.

Currently S3D content creators are reluctant to introduce negative parallax, and instead much of depth in the scene is beyond the plane of the screen. In this case, again assuming the left eye sees the original content, an object must be shifted to the right in the right eye, creating a monocular region on the left side of the object. This region has now been revealed to the viewer, and is not situated in an occluded zone (according to the scene geometry) so it should be visible to both eyes, but it is not. If not textured appropriately monocular gaps in this arrangement will cause rivalry and reduced depth perception [12, 11]. Consequently, this region must be carefully rendered, for inconsistencies in texture or luminance will draw the viewer's attention to this region. In arrangements geometrically consistent with occlusion the visual system will tolerate small texture variations within the monocular region [12]. However, the amount of perceived depth and the perceptual latency in such situations might be compromised [13, 14]. Moreover, if a scene contains several disjoint monocular regions, such as in a background viewed behind a fence, the monocular texture must be continuous from one region to another in order to ensure veridical depth perception [16].

Monocular occlusions and depth perception latency

The presence of consistent monocular occlusion regions can speed up processing of depth from disparity in both simple stereograms [10] and orthostereoscopic photographs of real objects [18]. In both cases, observers took less time to determine if the depth ordering in a scene was correct when self-occluded regions were present. It seems likely that monocular occluded aid in the rapid identification of depth discontinuities (see also [19]).

Relevance to S3D media

These findings emphasize the importance of correct rendering of monocular areas in S3D content. Correctly rendered monocular areas will enhance the ease with which depth is perceived in S3D media. It is possible (though not yet proven) that the appropriate use of monocular occlusion information may help offset some of the conflicting signals present in S3D displays. For instance, under near viewing conditions where an audience is positioned say 1 m from a display, the conflict between accommodation and vergence has been shown to cause eyestrain [20] . It is likely that manipulations which make the S3D more closely approximate vision in the natural environment will help ease the burden on the visual system and make the experience more comfortable.

Illusory contour and surface perception from monocular occlusions

In Lawson and Gulick's [9] early experiments they showed that monocular elements added along the boundary of a zero-disparity dot-defined rectangle causes the central blank region to be displaced in depth. What is often overlooked in their work is the fact that their stimuli consistently evoked a percept of a rectangular surface displaced in depth (see also [8] and Figure 3 above). Lawson and Gulick's sparse configuration could have promoted any number of percepts, including one in which the displaced outline was an empty wire frame. Instead the blank region consistently appeared opaque or solid. Similar, stable, illusory surface percepts from monocular occlusions have been reported in a number of studies [21, 8, 22, 23] . In sum, it appears that when possible the visual system will attempt to consolidate a monocular region into a scene by creating the percept of occluding surfaces and boundaries to account for the absence of a feature in one eye's view. This phenomenon has lead to the proposal that monocular occlusions serve an important role in defining surface boundaries, and in doing so provides a constraint for surface smoothing operations [24]. Nakayama and Shimojo [11] demonstrated that this is true, but the effect was weak in their stimulus. In a more recent study Tsirlin, Wilcox and Allison [25] have shown that, given appropriate geometric constraints, a monocularly occluded region can dramatically influence the perceived shape of an illusory surface (Figure 5). Moreover, illusory occluders can also capture binocularly defined features with ambiguous disparity such as repetitive patterns [19, 26].

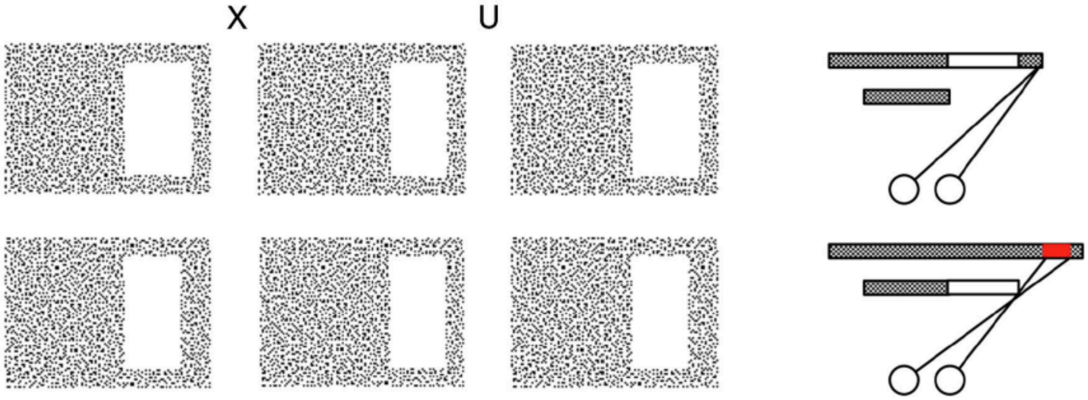

Figure 5. Monocular occlusions generate illusory occluding surfaces. Random-dot stereograms (rds) in the left and center columns are arranged for crossed fusion and in the center and right columns for uncrossed fusion. The rightmost column illustrates the typical corresponding cyclopean percepts where the random-dot texture is shown in gray, the blank region in white and occluded regions in red. In the top row there is no monocular region to the right of the blank area and it is perceived as part of the background. In the bottom row there is an occluded region of dots in the right-eye's image to the right of the blank area and now it is perceived as part of the foreground (figure modified from Tsirlin et al. [25]).

Relevance to S3D media

Accidental monocular occlusions that occur due to scratching, glare and other artifacts, can be interpreted by the visual system as monocularly occluded features. Depending on the configuration of the scene these features could either cause rivalry or be interpreted as indicating a depth discontinuity. This could have a disruptive effect on the organization of depth in a scene and draw viewers' attention to an arbitrary region of the scene. On the other hand, S3D content creators could also exploit the illusory occluder phenomenon. For example CGI animators could generate depth percepts in monochromatic surfaces which otherwise lack texture or disparity information.

Conclusions

The visual system is designed to make use of monocular occlusions to identify depth discontinuities, resolve ambiguous disparities, and in some cases define surface shape. Thus, the presence of monocular regions that do not correspond to the binocular geometry of the scene will cause rivalry and potentially disrupt depth percepts in S3D content. Here we have discussed two primary types of monocular occlusion: background and self-occlusion. We have classified monocular occlusions in S3D media as resulting from one of three different sources where they are added inadvertently, or via geometric distortion, or when the occluded region must be simulated. The perceptual effects of introducing monocular regions where they should not exist, or filling monocular content that is inappropriate for the particular scene, are significant. As we have outlined here, these consequences must be carefully considered when capturing, modifying or displaying S3D content.

Bibliography

1. I.P. Howard and B.J. Rogers, *Seeing In Depth*. Vol. 2. 2002, Toronto: Porteous Press.
2. L. da Vinci, *Selections from the notebooks of Leonardo da Vinci*, ed. R. I.A. 1977, Oxford: Oxford University Press.
3. J.M. Harris and L.M. Wilcox, "The role of monocularly visible regions in depth and surface perception", Vision Research, 49: 2666-2685, 2009.
4. T.L. Ooi and Z.J. He, "Binocular rivalry and surface-boundary processing", Perception, 35: 581-603, 2006.
5. C. Wheatstone, "Contributions to the physiology of vision. Part the first: On some remarkable and hitherto unobserved phenomena of binocular vision.", Philosophical Transactions of the Royal Society of London II: 371-394, 1838.
6. D. Marr and T. Poggio, "Cooperative Computation of Stereo Disparity", Science, 194: 283-287, 1976.
7. D. Marr and T. Poggio, "A computational theory of human stereo vision", Proc. R. Soc. Lond. B., 204: 301-328, 1979.
8. W.H. Ehrenstein and B.J. Gillam, "Early demonstrations of subjective contours, amodal completion, and depth from half-occlusions: "Stereoscopic experiments with silhouettes" by Adolf von Szily (1921)", Perception, 27: 1407-1416, 1998.
9. R.B. Lawson and W.L. Gulick, "Stereopsis and anomalous contour", Vision Research, 7: 271-297, 1967.
10. B. Gillam and E. Borsting, "The role of monocular regions in stereoscopic displays", Perception, 17: 603-608, 1988.
11. K. Nakayama and S. Shimojo, "Da Vinci stereopsis: Depth and subjective occluding contours from unpaired image points. Special Issue: Optics, physiology and vision", Vision Research, 30: 1811-1825, 1990.
12. S. Shimojo and N. K., "Real world occlusion constraints and binocular rivalry", Vision Research, 30: 69-80, 1990.
13. P.M. Grove and H. Ono, "Ecologically invalid monocular texture leads to longer perceptual latencies in random-dot stereograms", Perception, 28: 627-639, 1999.
14. P.M. Grove, et al., "Content and context of monocular regions determine perceived depth in random dot, unpaired background and phantom stereograms", Vision Research, 42: 1859-1870, 2002.
15. E.A. van Bogaert, et al., "The monocular-boundary-contour mechanism in binocular surface representation and suppression", Perception, 37: 1197-1215, 2008.
16. J. Forte, et al., "Binocular integration of partially occluded surfaces", Vision Research, 42: 1225-1235, 2002.
17. I.P. Howard, "Depth from binocular rivalry without spatial disparity.", Perception, 24: 67-74, 1995.
18. L.M. Wilcox and D.C. Lakra, "Depth from binocular half-occlusions in stereoscopic images of natural scenes", Perception, 36: 830-839, 2007.

19. B.L. Anderson and K. Nakayama, "Toward a general theory of stereopsis: Binocular matching, occluding contours, and fusion", Psychological Review, 101: 414-445, 1994.

20. D.M. Hoffman, et al., "Vergence-accommodation conflicts hinder visual performance and cause visual fatigue", Journal of Vision, 8, 2008.

21. B.L. Anderson, "The role of partial occlusion in stereopsis", Nature, 367: 365-368, 1994.

22. B. Gillam and K. Nakayama, "Quantitative depth for a phantom surface can be based on cyclopean occlusion cues alone", Vision Research, 39: 109-112, 1999.

23. L. Liu, et al., "Quantitative stereoscopic depth without binocular correspondence", Nature, 367: 66-69, 1994.

24. W.E.L. Grimson, *From Images to Surfaces*. MIT Press Series in Artificial Intelligence. 1981, Cambridge, MA: MIT Press. 274.

25. I. Tsirlin, et al., "Monocular occlusions determine the perceived shape and depth of surfaces", Journal of Vision, in press, 2010.

26. J. Hakkinen and G. Nyman, "Phantom surface captures stereopsis", Vision Research, 41: 187-199, 2001.