Displays 48 (2017) 35–40

Contents lists available at ScienceDirect

Displays

journal homepage: www.elsevier.com/locate/displa

True stereoscopic 3D cannot be simulated by shifting 2D content off the screen plane

Paul Hands, Jenny C.A. Read *

Institute of Neuroscience, Newcastle University, Tyne and Wear, United Kingdom

A R T I C L E I N F O

Article history:
Received 20 September 2016
Received in revised form 26 December 2016
Accepted 22 February 2017
Available online 1 March 2017

A B S T R A C T

Generating stereoscopic 3D (S3D) content is expensive, so industry producers sometimes attempt to save money by including brief sections of 2D content displayed with a uniform disparity, i.e. the 2D image is geometrically shifted behind the screen plane. This manipulation is believed to produce an illusion of depth which, while not as powerful as true S3D, is nevertheless more compelling than simple 2D. Our study examined whether this belief is correct. 30 s clips from a nature documentary were shown in the original S3D, in ordinary 2D and in shifted versions of S3D and 2D. Participants were asked to determine the impression of depth on a 7 point Likert scale. There was a clear and highly significant difference between the S3D depth perception (mean 6.03) and the shifted 2D depth perception (mean 4.13) ($P = 0.002$, ANOVA). There was no difference between ordinary 2D presented on the screen plane, and the shifted 2D. We conclude that the shifted 2D method not only fails to mimic the depth effect of true S3D, it in fact has no benefit over ordinary 2D in terms of the depth illusion created. This could impact viewing habits of people who notice the difference in depth quality.

1. Introduction

Leonardo da Vinci famously complained that flat paintings could never give a true impression of depth, because in real scenes the two eyes see different aspects of an object [1]. Since Wheatstone's [2] invention of the stereoscope, this limitation has been overcome, and today many forms of technology exist which are able to show the left and right eye a slightly different image of the same scene, including polarised light filters, active shutter glasses and parallax barriers [3–7]. Advances in digital technology mean that stereoscopic 3D (S3D) displays are more accessible than ever before. Consumers are now able to possess 3D-capable television sets in their own home [8,9]; several videogames manufacturers have produced 3D versions [10] and a number of companies are developing virtual reality headsets which incorporate S3D [11].

However, S3D content, especially live-action, remains complex and expensive to produce. A production standard mirror rig setup (including cameras) for S3D filming can easily cost more than $1,000,000. Given that filming an event usually requires many different camera angles and hence many different rigs, filming a football game in S3D could require as much as $10,000,000 of equipment (based on a minimum of 9 cameras needed, although typically the average is 12–15). These rigs have to be very precisely aligned to avoid distortions, and usually require extra personnel to operate, e.g. specialist 3D focus/convergence/interocular pullers as well as stereographers. Extra consideration also needs to be given to editing, since when changing aspects such as color saturation and brightness, both eyes need to be changed equally or distortions quickly appear [12,13]. Even for computer-generated S3D content, more rendering hours and more calculations are needed. Sometimes more than two renderings of the same scene are required, since the stereographer may decide that different regions of the scene need to be rendered with different camera parameters.

Given these issues, producers of S3D content occasionally use a shortcut rather than capturing every scene in S3D or converting it to S3D in post-production. They take 2D content and simply replicate the single camera lens image in both the left and the right eye, after offsetting them horizontally in opposite directions. The effect of this, geometrically, is to shift the planar 2D image back behind the screen plane; accordingly, we will call this "shifted 2D". The shift has to be behind rather than in front of the screen plane to avoid window violations. This shift is believed in the industry to create an illusion of depth which, while not as compelling as true S3D, is nevertheless more impressive than conventional 2D.

This belief is not unreasonable. 2D images contain many pictorial cues to three-dimensional structure, including perspective, shading, texture cues and apparent size. These can even trigger reflex vergence eye movements, implying that the brain accepts

* Corresponding author.
E-mail addresses: paul.hands@ncl.ac.uk (P. Hands), jenny.read@ncl.ac.uk (J.C.A. Read).

http://dx.doi.org/10.1016/j.displa.2017.02.002
0141-9382/

these depth cues at a basic perceptual, rather than simply cognitive, level. 2D video content includes still more powerful depth cues, such as structure from motion [14]. However, there is evidence that the visual system detects the flat picture plane, and that perception is powerfully influenced by this. Indeed, this seems to be a key reason why pictures and photographs look "correct" across a wide range of viewing angles, even though the image on the retina is changing profoundly. The visual system appears to detect the screen plane and correct for the oblique viewing angle. An undesirable side-effect of this is that we remain aware at some level that the image is "only a picture", projected onto a flat screen plane rather than genuinely existing in three-dimensional space.

There is a wealth of evidence, going back at least to Tscherning [15], that weakening the cues to the existence of the screen plane results in a stronger impression of depth. Binocular disparity is a powerful cue to the flatness of the screen plane, so weakening disparity cues is an immediate way of reducing the salience of the screen plane. Tscherning discusses the depth illusion produced when 2D pictures are viewed through Javal's iconoscope, an optical device which presents the same image to both eyes. Simply viewing a picture from a greater distance produces a similar effect, but the iconoscope also disrupts the relationship between convergence and viewing distance, a manipulation which itself increases the depth illusion [16]. The zograscope [17] worked in a similar way. Claparède [18] discussed the "paradox of monocular stereopsis": the stronger depth illusion created when 2D pictures are viewed monocularly, again because this removes the binocular cues to flatness. Ames [16] reports that blurring the image in one eye also strengthens the depth illusion, especially if a cylindrical lens is used to blur vertical lines while leaving horizontal ones sharp. Again, this presumably disrupts disparity cues to flatness, leaving pictorial cues free to dominate. Binocular cues are not the only ones indicating the screen plane. Accommodation is a monocular cue to flatness, at least at near viewing distances, so removing this cue (by viewing through a small hole) or disrupting it (by viewing through positive or negative lenses) also strengthens the illusion of depth. Ames reports that viewing a flat image through a mirror produces the same effect. This is presumably by introducing uncertainty as to the position of the picture in space: the frame removes the continuity between the observer and the picture via the surrounding objects and surfaces, while the mirror's surface presents a competing candidate for picture plane. Perhaps most interestingly for the present paper, Ames also discusses "changing the convergence of the eyes from that normally required by the distance from which the picture is viewed", by placing prisms in front of the eyes [19]. This is directly equivalent to the "shifted 2D" exploited by current S3D producers.

Thus, there are good grounds for expecting "shifted 2D" to produce a stronger illusion of depth than "native 2D" presented on the screen plane. We expect the shift to stimulate a vergence movement such that viewers no longer fixate the physical screen, but verge behind it, at the plane indicated by the parallax shift. As in Ames' mirror experiment, the physical screen is now geometrically a pane of glass through which the picture content is viewed, as through a window. The suggestion is that this may reduce the conflict between the flatness of the physical screen plane and the depth structure indicated by monocular depth cues within the content, allowing a stronger impression of depth. However, to our knowledge this suggestion has not yet been tested. The early literature was purely qualitative, and although more recently many of these effects have been examined quantitatively [17,20,21], as far as we are aware the present study represents the first quantitative examination of the effect of the "shifted 2D" manipulation on the experience of depth.

2. Material and methods

2.1. Equipment

The stimuli were displayed on a passive stereoscopic 3D display monitor (AOC D2367ph, http://www.aocmonitorap.com/v2015/nz/product_display.php?id=409) in a room which had regular, constant background luminance of 161.2 cd/m^2 (average of ten measurements made using a Minolta LS-100 photometer). The monitor resolution was 1920 × 1080 pixels, 47.6 cm wide × 26.8 high (diagonal 54.6 cm or 21.5″). The monitor was of the patterned-retarder type where left and right images are separated by circular polarisation and displayed on alternate pixel rows, halving the vertical resolution.

A chinrest was used to ensure that each subject viewed the content from the same position both horizontally and vertically with each trial, to ensure other effects, such as viewing distance and viewing angle [22,23], were not factors in determining immersion. The viewing distance was 100 cm. A height-adjustable chair was used to ensure the participants were comfortable during the experiment. Participants wore passive S3D glasses throughout the experiment. They were not told anything about whether the content would be in 2D or S3D.

2.2. Stimuli

The stimuli were 13 separate 30 s clips from the BSkyB production 'Micro Monsters with Sir David Attenborough', which was filmed in S3D. Clips were chosen from 2 episodes that were made available by BSkyB for the study, and were chosen so that the 30 s timespan started and ended at a sensible place, avoiding starting or stopping the clip midsentence, and also to be sure the clips were engaging. Both the left and right eye of each clip was also made available in an AVI file for the study. The software program 'Stereo Movie Maker' (available at http://stereo.jpn.org/eng/stvmkr/index.html) was used to modify the clips. The subsequent modified clips were each displayed in four different ways:

- Native S3D – showing the left clip to the left eye and the right clip to the right eye, as typically done in S3D content displays.
- Native 2D – showing the left clip to both the left and right eye. (Note that this will have been different to the broadcast 2D version, as a different editing procedure will have been used for the 2D footage.)
- Shifted S3D – as for Native S3D but in this case the left image was shifted left by 56 pixels and the right image was shifted right 56 pixels.
- Shifted 2D – as for Native 2D but this time the original left clip was shifted left by 56 pixels and displayed to the left eye, while the *same* clip was shifted right by 56 pixels and displayed to the right eye.

The shifted 2D condition contains the basic manipulation being examined in this paper. We wish to see whether it is true that it produces a better depth than the native 2D condition, and whether it approaches the depth quality of the native S3D condition. The shifted S3D condition enables us to probe the effect of the shift separately from the effect of S3D.

2.3. Applying the shift

The Stereo Movie Maker was used to create both the native clips and the shifted ones. For the native clips, both left and right half-images, with a resolution of 1920 × 1080 pixels, are stored in a side by side format and downsampled to 2048 × 576 pixels. When displayed on the 3D monitor in its 3D side-by-side mode, each eye's image occupies the full width of the screen but only half its height, i.e. extends over 47.6 cm wide × 13.4 cm high. The aspect ratio of the image is therefore changed from the original 1.78 (1920:1080) to 3.56 (2048:576).

To make the shifted clips, the leftmost 56 columns of pixels are removed from the left image and the rightmost 56 columns from the right image, resulting in half-images with a resolution of 1864 × 1080 pixels. Again the Stereo Movie Maker reduces the resolution, storing the clip at a resolution of 2048 × 592 pixels. When displayed on the 3D monitor in its 3D side-by-side mode, each eye's image extends over 47.6 cm wide × 13.8 cm high for an aspect ratio of 3.46 (2048:592). Note that the aspect ratio is slightly different for the shifted clips, because columns of pixels have been removed as part of the shift. When the image-pair is shrunk down to occupy 2048 pixels horizontally, it therefore contains slightly more pixels vertically (see Fig. 1)

Mathematically, pixels m_L in the original left image now map onto pixels n_L in the new image, where $n_L = (m_L - 57)1919/1863 + 1$, the valid values of n_L are 1 to 1920 and the valid values of m_L 57 to 1920 (the leftmost 56 pixels having been removed). In the right image, $n_R = (m_R - 1)1919/1863 + 1$, where the valid values of n_R are again 1 to 1920 but the valid values of m_R are 1 to 1863 (the right most 56 pixels having been removed). The resulting on-screen parallax introduced between corresponding pixels ($m_L = m_R$) is $(n_R - n_L) = 56 \times 1919/1863 = 57.7$ pixels or $P = 1.43$ cm. Geometrically, this is equivalent to displaying flat 2D content on a plane behind the monitor screen (Fig. 2), as if viewed through a glass window. Applying a parallax P in this way increases the geometrically-defined distance by a factor $I/(I - P)$, where I is the observer's interocular distance, in the same units as P. This factor is independent of the screen width and of the viewing distance. For our experiment, the parallax was 1.43 cm, meaning that for an observer with eyes 6.5 cm apart, the geometry specified an increase in distance by a factor of 1.28. Thus at our 100 cm-viewing distance, the shift places the virtual content 128 cm from the observer according to the binocular geometry (Fig. 2).

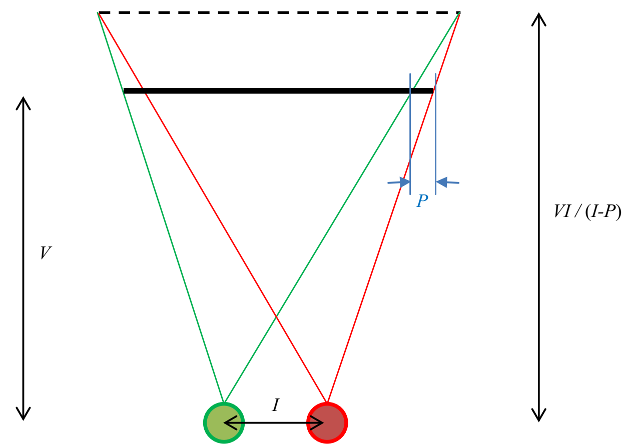

Fig. 2. Geometry of our experiment. The interocular distance I is for calculations assumed to be 6.5 cm. By applying a parallax of 1.43 cm, the virtual image is moved behind the screen by a factor of 1.28. So when the viewing distance V is 100 cm, the plane of the image should appear to be 128 cm away.

Fig. 1. (a) Right and left images displayed for crossfusion for the native S3D (top row) and shifted 2D (bottom row). (b) The shifted 2D image as it appeared on the row-interleaved screen.

This resulted in 52 different stimuli for the participant to look at, 4 versions of each of 13 clips. The 52 stimuli were shown in the same order to each of 9 participants. The four versions of each clip were presented consecutively, with the four conditions coming in a different order for each clip.

Due to a formatting issue with creating the shifted clips, the frame rates were 25 fps for the native and 12 fps for the shifted clips. Participants did not report noticing any differences between the frame rate of the clips in terms of quality or judder, and indeed the authors could not reliably detect which clips were displayed at which frame rate. As shown below, the results did not correlate with frame rate.

2.4. Procedure

Participants were asked to sit in the chair comfortably, wearing the glasses and resting their chins on the chinrest. Eye movements were tracked using an Eyelink 1000 eyetracker (SR Research Ltd, http://www.sr-research.com/EL_1000.html), on an angled binocular configuration with a 25 mm wheel lens. The eyetracker was positioned 55 cm away from the participant, underneath the monitor in their visual field. Once they were comfortably positioned the participant viewed all the clips. After each clip the participants were asked to assess the perceived depth in the image, stressing that the actual content (i.e. how interesting it was) wasn't important to the study, and to give a score on a 7 point Likert scale, from 1 being "no noticeable 3D" to 7 being "fantastic, immersive 3D". Participants were given a £10 shopping voucher for their participation.

2.5. Participants

Participants were recruited via an internal volunteer scheme at Newcastle University Institute of Neuroscience, on the basis that they had no visual problems. The work was approved by Newcastle University Faculty of Medical Sciences Ethics Committee. 9 participants (6 female, all naïve; 3 male, author PH and 2 naïve) were used in the study. Naïve participants were not informed of the experimental aims or hypotheses.

3. Results

3.1. Vergence

Fig. 3 shows the vergence for the 9 participants in the four different conditions. This is plotted as parallax on the screen plane, i.e. the horizontal distance between the pixel viewed by the left eye and that viewed by the right eye. The error on this measurement for each eye is comparable to the parallax, so these data are extremely noisy. There are also systematic differences between participants which may reflect individuals' fixation disparity or errors in calibration. Nevertheless, after averaging across clips and participants, we do see a clear effect of shift. On average, participants converge on the screen plane in the "native" conditions (a *t*-test on mean parallax for each participant revealed no significant difference from zero in either of the native conditions). This is as expected: native 2D content is of course exactly on the screen plane; native S3D content may be behind or in front of the screen, but is close to the screen plane on average. In the "shifted" conditions, participants converge behind the screen plane, close to the plane of the content (*t*-tests indicate a significant difference from zero parallax, $p < 0.01$ for shifted 2D, $p = 0.01$ for shifted S3D; *t*-tests indicate no significant difference from shift parallax 1.4 cm). A two-way Analysis of Variance with "Shift" (shifted vs native) and "Stereo" (2D vs S3D) as factors indicated a significant effect of Shift ($p < 0.01$, $F = 7.52$, $df = 1$) but no effect of Stereo and no interaction. This confirms that our shift condition did induce the expected change in vergence.

3.2. Depth quality ratings

Fig. 4 shows the average score for each different viewing condition, (a) for the different subjects and (b) for the different clips. It is immediately clear that the depth ratings are substantially higher for S3D (red, triangles) than for 2D (blue, squares). However, there is very little difference between the native and shifted formats. For S3D, shifting has no effect (mean rating 6.02 for native vs 6.03 for shifted). For 2D, depth ratings are marginally higher for shifted (mean rating 4.03 for native vs 4.23 for shifted), but this difference is not significant. A two-way repeated-measures ANOVA on each subject's average ratings across the 13 clips, with stereo (2D vs S3D) and plane (native vs shifted) as factors, found a highly significant main effect of stereo condition ($F = 19.9$, $P = 0.002$), but no main effect of plane ($F = 1.10$, $P = 0.33$) and no interaction between plane and stereo condition ($F = 0.634$, $P = 0.45$). Since the shifted conditions were scored the same as the corresponding native conditions, we can conclude that neither the shift itself, nor the confounded change in aspect ratio and frame rate, had a significant effect on depth quality ratings. There was no evidence that the clips themselves differed systematically in the depth impression they produced. For example, a Kruskal-Wallis test finds no difference between the ratings given to the 13 different native 2D clips ($P = 0.44$).

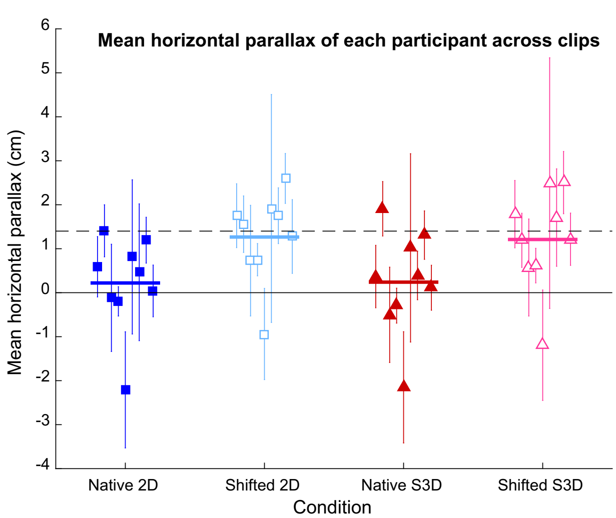

Fig. 3. Mean horizontal screen parallax of each participant's gaze, averaged over time for all clips in a given condition. The eyetracker estimated where on the screen the left and right eyes are looking (x_L, y_L) and (x_R, y_R); horizontal screen parallax is the difference in the x-coordinates. The eyetracker reports these in pixels but we have converted to cm for display purposes. Errorbars show ± 1 standard deviation (standard errors on the mean are tiny because thousands of samples were taken). The black lines indicate the screen parallax of the content, i.e. zero for the native conditions (solid) and 1.4 cm for the shifted (dashed). Colored lines show the mean across participants for each condition. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4. Discussion

The results confirm previous results showing, unsurprisingly, that viewers experience a more impressive illusion of depth with S3D as compared to 2D content. The results are very much in line with Bohr and Read [24], who also used a 7-point Likert scale to investigate depth realism (their Fig. 6), this time across different groups who viewed the film "Toy Story" in either 2D or S3D [24]. The mean "depth realism" rating was 5.40 for their two S3D groups, compared to 4.26 for their three 2D groups ($P < 10^{-10}$, Mann-Whitney rank sum test; there were no significant differences within their S3D or 2D groups). The slightly smaller difference between 2D and S3D in Bohr & Read [24] may reflect their between-subjects comparison; no participants had the opportunity to compare 2D and S3D directly as in the present study, where all participants viewed all clips in all conditions.

We found no evidence that shifting 2D content behind the screen plane produces a stronger illusion of depth. Depth ratings were very slightly higher for shifted 2D content, and this difference might possibly have become significant if we had more statistical power. However, for practical purposes this is immaterial. The question at issue was whether the depth shift could simulate the depth of true stereoscopic 3D, and here the answer is clear: it does not come close. Even if the increase could be shown to be significant by a more powerful study, it would still be too small to be of interest as a practical way of substituting for true S3D. Apparently, the binocular disparity cues indicating that the content is flat still dominate, even when the image is shifted behind the physical screen plane.

This conclusion is necessarily limited to the particular clips we used. These were all taken from the same S3D film "Micro Monsters", they were all similar in nature (wildlife documentary), and they did not differ in the strength of the depth illusion they created. The logic of the shift manipulation is that weakening cues to the screen plane enables monocular depth cues to dominate perception [16]. This would predict that the effect of the shift should be stronger for content with more powerful monocular depth cues. As Koenderink et al. [17] write, "a photograph of a brick wall in frontoparallel attitude is not going to reveal any 'zograscopic effect' [illusory depth]". More subtly, a 2D photograph of several frontoparallel surfaces at different distances may also display very little zograscopic effect, simply because the monocular depth cues are weak, even though the binocular disparity cues make the surfaces appear clearly separated in depth when viewed in S3D. Could this

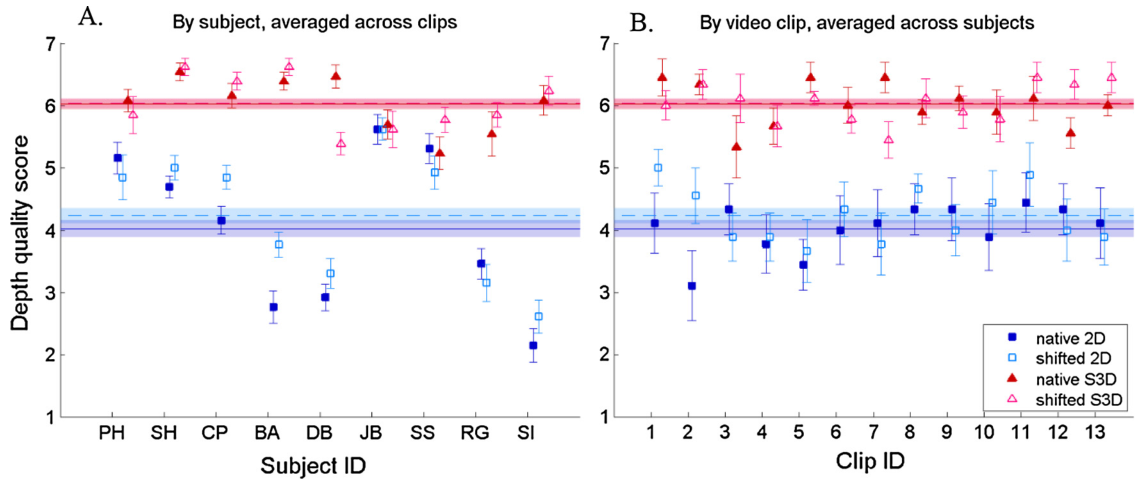

Fig. 4. Depth quality scores for each of the four different viewing conditions, (a) for the 9 different subjects, averaged across the 13 video clips, and (b) for the 13 different clips, averaged across subjects. Blue squares show results for 2D, red triangles for S3D; filled symbols/solid lines are for native content, empty symbols/dashed lines are for content shifted behind the screen plane. Errorbars show ±1SEM of the 9 subjects' judgments for each data-point; points are offset horizontally so that errorbars do not overlap. Horizontal lines show means for each condition, averaged across content & subjects. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

explain our results? The highly significant increase in depth ratings when the content was displayed in S3D proves that, unlike the brick wall example, our scenes did depict a wide range of depth, but we have not assessed their monocular depth cues objectively. The literature on "monocular stereopsis" to date contains only cursory discussion of how the nature of the content might affect the strength of the effect [17]. We cannot identify any particular reason why the content we used should be particularly ineffective at producing a "zograscopic effect". The literature reviewed in the Introduction used static images, which cannot contain depth cues such as structure-from-motion and looming, whereas our "Micro Monsters" clips regularly contained these cues. Additionally, the clips were typically of insects filmed in extreme macro, meaning that they contained depth-of-field (blur) cues to three-dimensional structure. Rather than consisting of sets of frontoparallel surfaces with little depth structure within each surface, the clips typically depicted undergrowth, bark and so on extending in depth. Thus, while we cannot rule out that the shift manipulation would have produced a more compelling depth impression with other content, our chosen examples seem likely to have had monocular depth cues at least as strong as other commercial S3D content, meaning that they should have been more potent in creating a depth illusion once the stereoscopic cues to flatness were weakened.

The shift we applied may simply have been too small to produce the intended effect. Our images were 48 cm across and shifted so as to give a screen parallax of 1.4 cm. This means that the binocular geometry specified the content as being at a viewing distance of 128 cm, 28 cm behind the physical screen plane at 100 cm. This is a substantial parallax, representing 3% of the image width (1.43 cm out of 47.6 cm). BSkyB's Technical Guidelines for Plano Stereoscopic (3D) Programme Content (http://www.sky.com/shop/__PDF/3D/3D_Tech_Spec_Short_Rev5_CJ_Apr2011.pdf) specify that parallax behind the screen "should not exceed 2% for majority of shots" (the limit for parallax in front of the screen is even smaller at 1%). Thus, the parallax we applied is substantial by the standards of commercial S3D content and it is unlikely to be practical to apply larger amounts. However, Ames [16] recommends 3Δ prisms (i.e. prisms which rotate the images through 1.72°) for images viewed at a distance of 100 cm. To achieve this disparity, we would have to shift each eye's image 3 cm on the screen, for a total parallax of 6 cm or 13%. Thus, the most likely reason for our failure to see a "zograscopic effect" is simply that much larger disparities are required. If this is the case, it means that the technique is not suitable for use in broadcast content. A screen parallax of 6 cm at a viewing distance of 100 cm would represent a vergence/accommodation conflict of around 1 dioptre, well outside the zone of comfort [25]. On displays larger than ~50 cm in width, a prallax of 13% would cause ocular divergence, which is extremely uncomfortable or even painful.

It is possible that more subtle factors than just the limited parallax account for our failure to find a depth enhancement. Prisms, as used by Ames, probably also distort the image and create much greater uncertainty about the location of the picture in space than in our set-up, where we retain all the usual cues indicating that the monitor is physically 100 cm in front of the observer. More subtly, prisms have a different effect on vertical disparity. Our manipulation introduces purely horizontal screen parallax

In summary, we have found no evidence that shifting 2D content behind the screen produces a depth illusion that is at all comparable to true S3D, at least not without the use of unacceptably large parallaxes. We conclude, regrettably, that the technique is not viable as a cheap way of making "fake" S3D.

Acknowledgements

Funding: This work was supported by BSkyB (SKY-LDM. FID58255, Grant Way, TW7 5QD) and EPSRC (EP/J500288/1) as a CASE PhD studentship supporting PH. Thanks to Ann Fitchett for running the Institute of Neuroscience volunteer recruitment scheme. Thanks to Chris Johns and staff at BSkyB for assistance preparing the film clips. All data presented here was collected by author PH for use in this study.

References

[1] L. da Vinci, Treatise on Painting (transl. J. F. Rigaud), second ed., J.B. Nichols & Son, London, 1835.
[2] C. Wheatstone, On some remarkable, and hitherto unobserved, phenomena of binocular vision, Philos. Trans. R. Soc. Lond. 128 (1838) 371–394.
[3] R. Snowden, P. Thompson, T. Troscianko, Basic Vision: An Introduction to Visual Perception, second ed., Oxford University Press, Oxford, 2012.
[4] M.S. Banks, J.C. Read, R.S. Allison, S.J. Watt, Stereoscopy and the human visual system, SMPTE Motion Imaging J. 121 (2012) 24–43, http://dx.doi.org/10.5594/j18173.

[5] S. Pastoor, M. Wöpking, 3-D displays: a review of current technologies, Displays 17 (1997) 100–110, http://dx.doi.org/10.1016/S0141-9382(96)01040-2.
[6] R. Burks, C. Harper, M.C. Bartha, Examining 3-D technologies in laptop displays, Ergon. Des. Q. Hum. Factors Appl. 22 (2014) 17–22, http://dx.doi.org/10.1177/1064804614526197.
[7] F. Devernay, P. Beardsley, Stereoscopic cinema, Image Geom. Process. 3-D Cinematogr. 35 (2010) 11–51, http://dx.doi.org/10.1007/978-3-642-12392-4.
[8] J.C.A. Read, Viewer experience with stereoscopic 3D television in the home, Displays 35 (2014) 252–260.
[9] H. Karajeh, M. Maqableh, R. Masa'deh, A review on stereoscopic 3D: home entertainment for the twenty first century, 3D Res. 5 (2014) 26, http://dx.doi.org/10.1007/s13319-014-0026-3.
[10] J. Schild, M. Masuch, Entertainment Computing – ICEC 2011, Springer, Berlin, Heidelberg, 2011, http://dx.doi.org/10.1007/978-3-642-24500-8.
[11] K.S. Hale, K.M. Stanney (Eds.), Handbook of Virtual Environments: Design, Implementation, and Applications, second ed., CRC Press, Boca Raton, 2015.
[12] P. Cavanagh, Reconstructing the third dimension: Interactions between color, texture, motion, binocular disparity, and shape, Graph. Image Process. 37 (1987) 171–195, http://dx.doi.org/10.1016/S0734-189X(87)80001-4.
[13] A. Woods, T. Docherty, R. Koch, Image distortions in stereoscopic video systems, SPIE's Symp. Electron. Imaging Sci. Technol. 36 (1993) 36–48, http://dx.doi.org/10.1117/12.157041.
[14] D.L. Ringach, M.J. Hawken, R. Shapley, Binocular eye movements caused by the perception of three-dimensional structure from motion, Vision. Res. 36 (1996) 1479–1492, http://dx.doi.org/10.1016/0042-6989(95)00285-5.
[15] M. Tscherning, Physiologic Optics, second ed., The Keystone, Philadelphia, 1904.
[16] A. Ames, a. Ames Jr., The illusion of depth from single pictures, J. Opt. Soc. Am. 10 (1925) 137, http://dx.doi.org/10.1364/JOSA.10.000137.
[17] J. Koenderink, M. Wijntjes, A. van Doorn, Zograscopic viewing, Iperception 4 (2013) 192–206, http://dx.doi.org/10.1068/i0585.
[18] E. Claparède, Stéréoscopie monoculaire paradoxale, Ann. Ocul. 132 (1904) 465–466.
[19] A. Duane, The effect of converging prisms upon our notions of size and distance—an experimental study, Ophthal. Rec. 9 (1900) 595–607.
[20] J.J. Koenderink, A.J. van Doorn, A.M. Kappers, On so-called paradoxical monocular stereoscopy, Perception 23 (1994) 583–594. <http://www.ncbi.nlm.nih.gov/pubmed/7800471>.
[21] D. Vishwanath, P.B. Hibbard, Seeing in 3-d with just one eye: stereopsis without binocular vision, Psychol. Sci. 24 (2013) 1673–1685, http://dx.doi.org/10.1177/0956797613477867.
[22] E. Brenner, W.J.M. Van Damme, Perceived distance, shape and size, Vision. Res. 39 (1999) 975–986, http://dx.doi.org/10.1016/S0042-6989(98)00162-X.
[23] P. Hands, T.V. Smulders, J.C. Read, Stereoscopic 3-D content appears relatively veridical when viewed from an oblique angle, J. Vis. 15 (2015) 6, http://dx.doi.org/10.1167/15.5.6.
[24] J.C. Read, I. Bohr, User experience while viewing stereoscopic 3D television, Ergonomics 57 (2014) 1140–1153, http://dx.doi.org/10.1080/00140139.2014.914581.
[25] T. Shibata, J. Kim, D.M. Hoffman, M.S. Banks, The zone of comfort: predicting visual discomfort with stereo displays, J. Vis. 11 (2011), http://dx.doi.org/10.1167/11.8.11.

Paul Hands is a software developer at Scott Logic in Newcastle. He completed his Masters of Research in 2013 and received a PhD in Neuroscience in 2016 studying in Jenny Read's Lab, researching human perception in S3D compared to 2D. Before this he achieved a first class mathematics undergraduate degree at Newcastle University.

Jenny Read is Professor of Vision Science at Newcastle University's Institute of Neuroscience. She has a doctorate in astrophysics from Oxford University and did postdoctoral training at the US National Eye Institute. She now runs a multidisciplinary research lab examining many aspects of visual perception, especially 3D or stereoscopic vision. Current projects include modelling how visual cortex encodes binocular information, developing a new 3D vision test for children (http://research.ncl.ac.uk/asteroid/), a European Training Network on Full-parallax Imaging (http://www.full-parallax-imaging.eu/) and uncovering how insects see in 3D (http://www.jennyreadresearch.com/research/m3/).