

Stereo from Shading

Alexandre Chapiro^{1,2}, Carol O'Sullivan³, Wojciech Jarosz², Markus Gross^{1,2}, Aljoscha Smolic²,

¹ ETH Zurich, ² Disney Research Zurich, ³ Disney Research LA




Figure 1: Stereo from Shading: the anaglyph image on the right is assembled from a single view (center) with no parallax disparity, but with different shading (left), producing a 3D impression. The inset represents the difference between eye renders.

Abstract

We present a new method for creating and enhancing the stereoscopic 3D (S3D) sensation without using the parallax disparity between an image pair. S3D relies on a combination of cues to generate a feeling of depth, but only a few of these cues can easily be modified within a rendering pipeline without significantly changing the content. We explore one such cue—shading stereopsis—which to date has not been exploited for 3D rendering. By changing only the shading of objects between the left and right eye renders, we generate a noticeable increase in perceived depth. This effect can be used to create depth when applied to flat images, and to enhance depth when applied to shallow depth S3D images. Our method modifies the shading normals of objects or materials, such that it can be flexibly and selectively applied in complex scenes with arbitrary numbers and types of lights and indirect illumination. Our results show examples of rendered stills and video, as well as live action footage.

Categories and Subject Descriptors (according to ACM CCS): Hardware Architecture [I.3.1]: Three-dimensional displays—; Computer Graphics [I.3.3]: Three-Dimensional Graphics and Realism—Display Algorithms; Image Processing and Computer Vision [I.4.8]: Scene Analysis—Stereo;

1. Introduction

In natural perception, the human visual system combines several depth cues in order to create a sensation of depth. These include the binocular cues of convergence and disparity, as well as a variety of monocular cues, such as motion parallax, interposition, occlusion, size, perspective, accommodation, and others. Stereoscopic 3D (S3D) as enabled by stereo 3D displays supports binocular cues and from that creates a sensation of depth. Presenting a separate view to each eye does not, however, perfectly emulate real life depth perception so limitations such as the vergence-accommodation


conflict [SKHB11] and other imperfections remain problematic. Furthermore, all 3D displays are limited to a feasible depth volume. In particular, autostereoscopic displays are still very restricted in the depth range they can reproduce without artifacts due to limited angular resolution. Therefore, enhancement of S3D perception has been a very active field of recent research, which has been focused mainly on disparity enhancement [DRE*12b, CHA*14].

In this paper we explore the novel use of lighting variations between eye renders as a means to enhance S3D. We are inspired by Puerta [MP89], who showed that a 3D illu-

sion is created when images with no parallax but different cast shadows are displayed stereoscopically,

The goal of our new method, *shading stereo*, is to generate an increase in the S3D sensation by leveraging differences in shading between views, with cast shadows left unchanged between renders. Our algorithm does not change the position of light sources as in [MP89], but rather modifies the vertex normals of selected objects in a scene. This method, not previously used in 3D rendering, allows us to choose where in a scene to add shading stereo and to apply the method to complex scenes with multiple and varied light sources. Our main contributions are:

- The use of variable shading between eye renders as a tool to create or enhance S3D.
- An algorithm that applies shading stereo to scenes with arbitrary lighting by manipulating normals rather than lights.
- The application of shading stereo to live action scenes in scenarios where relighting is feasible.

All examples in the paper can be viewed with anaglyph glasses  (red-l, cyan-r). The supplemental images can be optimally viewed using color neutral glasses (such as time-multiplexed or polarized variants) if available, or alternatively anaglyph glasses.

2. Related Work

By displaying two photographs of a scene between which the light was shifted horizontally, Puerta [MP89] generated a 3D effect when viewing stereo images without disparity. They suggest that the effect is caused by the difference in cast shadows between views, but mention that shading could possibly also act as a stereo cue. Langer and Bülthoff [LB99] hypothesize that shading could be an effective cue to communicate the shape of objects, particularly under natural lighting conditions such as diffuse light or lighting from above. While differences in lighting have not previously been used to augment the depth perception of S3D, researchers have shown that shadows can increase speed and reduce error rates of depth ordering tasks [BRK03] or to improve visual processing of technical content [ŠPV11]. For a full discussion of the role of cast shadows in 3D perception, please see [KM14].

Recent research has focused on enhancing perceived depth by augmenting the disparity cue. Lang et al. [LHW*10] propose a mapping function that maps stereo into a target space non-linearly and Chapiro et al. [CHA*14] suggest a depth re-mapping approach to increase perceived depth under extreme compression. Masia et al. [MWA*13] present a content re-mapping approach for autostereo while avoiding blurriness induced from limited angular resolution. Didyk et al. [DRE*12a] take advantage of the Cornsweet illusion to create a convincing stereo experience and reduce the overall depth range. They also propose a perceived

disparity model that takes into account both contrast and disparity [DRE*12b]. All these methods target parallax disparity as the main source of depth perception and do not consider light and shadows. Some work has addressed the influence of color contrasts [IKA07] and, more recently, luminance differences [VMB*14] to S3D, but shadows and shading are not directly addressed. View-dependent effects like specular and refractive materials are particularly challenging when displaying stereoscopic content and have been addressed by several researchers [DKR*14, TDR*12]. Shading stereo could also be combined with these methods to enhance depth perception and handle highlights effectively.

In order to create stereoscopic content that is backwards compatible, Didyk et al. [DRE*11] present stereo techniques that compress disparity until it is barely noticeable without glasses, while maintaining a small noticeable depth effect. Others aim to create display technologies that show artifact-free content that can be viewed with or without glasses, while sacrificing some contrast [SLV*13]. Shading stereo could be used for this purpose either by itself or in combination with such approaches, since the lack of disparity and identical direct shadows between views means that the mixed view seen without glasses is only barely distinguishable from a regular monoscopic image.

Finally, with respect to perception and comfort in stereo, the vergence-accommodation problem [SKHB11], cardboarding [CDP*14] and motion and luminance effects [DMHG13] have all been investigated. The brain has been demonstrated to fuse different low dynamic range images presented to each eye to a higher dynamic range impression [YZWH12]. Changing the normals of objects in order to exaggerate details, or to help visually parse complex information [RBD06], has also been proposed, albeit not for S3D applications. In this paper, however, we present subtle lighting changes to each eye to create an S3D effect.

3. Shading Stereo

Figure 2 (top) shows a traditional S3D camera setup, with light positions static between views and two shifted cameras. Simply shifting light sources between views, as in [MP89], would cause the same object to cast shadows in different directions because of varying illumination (see Figure 3-A). Furthermore, since the lighting of the scene as a whole is changed between renders, some objects with complex geometries may acquire disturbingly different colors that are hard to fuse stereoscopically. Reflections and refractions may be affected, and multiple lights, or lights that are not conveniently located near the camera, would also be difficult to handle.

Our solution to this problem, illustrated in Figure 2 (bottom), consists of shifting the edits from the scene illumination to specific objects in the scene. In particular, since Lambert's cosine law states that the diffuse reflection off a point




Figure 2: Top: stereo from disparity. Bottom: shading stereo modifies normals to simulate shifting lights.




Figure 3: Left: rendering S3D with shifted lights. Right: shading stereo is applied to the teapot only.

is directly proportional to the cosine of the angle of incident light to the surface normal, we can directly edit the normals at all shade points of an object so that the illumination from a light positioned at 0 matches exactly that of a shifted light by employing an appropriate rotation. Now, arbitrary lighting of the scene can be employed, as the new lighting will affect the rest of the scene normally, while the object of shading stereo will have modified shading. If the lighting consists of a single point light in its original position, the new method on the edited object will match the shifted lighting described in the previous paragraphs exactly.

We begin by defining the camera position in 3D \mathbf{c} and two virtual shading points that are shifted along the camera axis symmetrically away from the camera, \mathbf{s}_l and \mathbf{s}_r . For each eye, we use one of these virtual points as appropriate, and the position of the virtual shading point being used for the current render is \mathbf{s} . Given a surface \mathcal{M} , for each point $\mathbf{x} \in \mathcal{M}$, we define $\vec{v}_c = \mathbf{c} - \mathbf{x}$ and $\vec{v}_s = \mathbf{s} - \mathbf{x}$ and perform the following operations:

- The axis of rotation is defined as $\vec{a}_x = \vec{v}_c \times \vec{v}_s$.
- If the original normal of \mathbf{x} is \vec{n}_x , the angle by which the normal is rotated is defined as $\theta_x = \theta_{xc} - \theta_{xs}$, where $\theta_{xc} = \arccos(\vec{n}_x \cdot \vec{v}_c)$ and $\theta_{xs} = \arccos(\vec{n}_x \cdot \vec{v}_s)$, i.e., how different the angle between the vertex normal and the camera is from the angle between the vertex normal and the virtual shading point. All vectors above are considered normalized, so a dot product is the cosine of the factors.




Figure 4: Shading stereo: the normal of \mathbf{x} is adjusted to match the angle of a virtual shading point.

- The new normal \vec{n}'_x is the result of rotating \vec{n}_x around the axis \vec{a}_x by an angle of θ_x .

This operation is performed for each point, thus obtaining the new normals. Figure 4 illustrates these operations. As for S3D rendering in general, the problem remains of handling specularities and reflections, which often produce incorrect depth sensations [TDR*12]. We render specularities separately and adjust them to what is defined as a “flat” setting by Templin et al. [TDR*12].

Finally, while shading stereo requires access to virtual lights and normals, it is possible to obtain S3D effects in live-action sequences through image relighting (see [vdPJD*14]). As a proof of concept, we demonstrate shading stereo on scenes where geometry was known a priori. We overlay the shaded geometry on the original camera views, thus generating a 3D sensation when viewed stereoscopically. The shading from an additional virtual light can thus be used to generate the shading stereo effect (see Figure 8). Alternatively, the lighting of a scene could be emulated in a rendering environment for a better match between shading and image.

4. Feasibility Study

Our method follows the premise that S3D sensation can be generated or increased by simply rendering the left and right views with different shading. However, only a small subset of all possible variations of illumination between both eyes will cause a plausible depth sensation. For example, we observed that the retinal rivalry caused by an exaggerated difference in shading generates discomfort. We therefore conducted a perceptual experiment to explore the feasibility of using shading stereo for S3D without causing disturbing retinal rivalry.

Stimuli: Scenes were presented on an Alienware 2310 23", 3D capable display with the help of time-multiplexed glasses. Participants were seated at a comfortable distance




Figure 5: Stimuli shown in the experiment, where $(dParallax, dLighting) =$ A: 0.0,0.0, B: 0.2,0.0, C: 0.0,0.4, D: 0.2,0.6, E: 0.6,0.6

of about 60 centimeters from the screen. The virtual camera setup attempted to emulate natural viewing by placing cameras on a scene 60 centimeters away from the stimuli and 6 centimeters away from each other, so that the virtual stimuli would have 3D characteristics of size and shape similar to a real world objects (orthostereo).

The rendering setup consisted of a stereo pair oriented in parallel along the depth axis, with the resulting images re-converged around the center of the stimuli, thereby displaying it with null parallax disparity. We consider this to be the unit baseline. Given a point with coordinates (x, y, z) and parallax disparity ψ , if the interaxial distance is reduced to $dParallax \in [0, 1]$, the disparity will change to be $dParallax \cdot \psi$. As such, depth can be controlled proportionally in our rendering system for experimental purposes, so $dParallax$ can be thought of as the depth compression factor.

Since the virtual shading points can be placed arbitrarily far from an object, an absolute distance such as that used for cameras would not be suitable for parametrization. We therefore parameterized the difference in shading by angle θ , the angle between the virtual shading points and the center. Given a maximum angle (in our experiments, $\pi/6$), we introduce a light disparity $dLighting \in [0, 1]$ so that the angle between the virtual shading points at any time is $dLighting \cdot \theta$.

Method: Fifteen naive volunteers (3F, 12M, ages 24–35, with normal or corrected-to-normal vision) participated in the experiment. We displayed simple geometric models as in [RFWB07], ranging from flat to bumpy (Fig. 5), at compression ratios of $dParallax \in \{0.0, 0.2, 0.4, 0.6\}$. We hypothesized that irregularities in the shape would generate more pronounced self-shading and hence seem more 3D. Users were given control of the parameter $dLighting$ using a method-of-adjustment procedure. They were tasked with setting the following three light disparity measures in turn: (i) the smallest $dLighting$ where they perceived a just noticeable difference in depth, compared to the presented $dParallax$ baseline; (ii) the smallest $dLighting$ where they first perceived retinal rivalry; and (iii) a preferred $dLighting$ setting between the two values found above.

We performed a $5(model) \times 4(dParallax) \times 3(measure)$ Repeated Measures Analysis of Variance (ANOVA) on partici-




Figure 6: Results of our feasibility study to determine acceptable perceptual limits for shading stereo. Error bars show standard errors.

pants' responses (see Fig. 6) and found a main effect of measure ($F(2, 28) = 156, p \approx 0$). Newman-Keuls post-hoc tests of differences between means showed that all three measures were significantly different ($p < 0.05$). We found no significant main or interaction effects of $dParallax$ or model. This simple experiment demonstrates that shading stereo does affect the perception of depth, within a limited range, but further studies are required to fully evaluate and quantify the effect.

5. Results and Discussion

We have explored the feasibility of using a lesser known stereo cue to create or enhance depth perception. Our algorithm is easy to implement and does not add significant complexity to the rendering process and can be applied to both images and video. The effect can be tuned and adjusted flexibly on a per-object or even per-material basis and can be used on scenes with arbitrary lighting.

Figure 7 reveals some interesting features of shading stereo. Firstly, the first and second columns show that shading stereo creates a noticeable depth illusion which is clearly visible when compared to the original 2D image. The third column shows a version of the content with no shading stereo, but with a small amount of parallax disparity ($dParallax = 0.2$). Note that the depth range of these examples corresponds approximately to the depth reproduction limits of current autostereoscopic displays. Finally, the third and fourth columns show that using shading stereo in combination with disparity enhances the depth illusion.

The last comparison above hints at interesting practical uses for shading stereo. While in general the depth range that can be generated by shading stereo without causing discomfort is significantly more limited than disparity, shading stereo can still be used to enhance the depth sensation on depth limited display devices (e.g. autostereoscopic displays) to go beyond the devices' capabilities. Furthermore, shading stereo images without disparity can be viewed without glasses and still provide reasonable 2D image quality. This enables backward compatible stereo applications, where the same image can be viewed in both 2D (without glasses) and 3D (with glasses) with good quality.

Our method has several limitations. Firstly, views have to be re-lit, which requires access to the renderer or image re-lighting methods. If geometric approximations of a scene are used, some degree of precision and alignment is required. Excessive use of light disparity causes retinal rivalry, and while we found some thresholds in this paper, complex scenes might have significant variation (see Fig. 9). A failure case where geometry is estimated poorly is also shown.

Finally, we have presented an initial exploration of shading as a factor for 3D perception, but many questions remain. Further studies are required to determine how shading stereo compares to disparity, what kind of lighting works best with shading stereo, which materials can and cannot be used, and other interesting directions for future work.




Figure 8: (L) geometry proxy re-lit using our method (left eye view) (R) anaglyph image showing stereo from shading. The insets represent the difference between left and right eye images for each example.

6. Acknowledgements

The geometry for the “pots” example was created by an artist. Geometries for the “face” example were kindly pro-

vided by Beeler et al. [BHB*11]. The geometry for the “couch” example was provided by Kim et al. [KZP*13].

References

- [BHB*11] BEELER T., HAHN F., BRADLEY D., BICKEL B., BEARDSLEY P., GOTSMAN C., SUMNER R. W., GROSS M.: High-quality passive facial performance capture using anchor frames. In *ACM Trans. Graph.* (2011), vol. 30, ACM, p. 75. [5](#)
- [BRK03] BAILEY M., REBOTIER T., KIRSH D.: Quantifying the relative roles of shadows, stereopsis, and focal accommodation in 3d visualization. In *IASTED VIIP* (2003), pp. 992–997. [2](#)
- [CDP*14] CHAPIRO A., DIAMANTI O., POULAKOS S., O’SULLIVAN C., SMOLIC A., GROSS M.: Perceptual evaluation of cardboarding in 3d content visualization. In *Proc. ACM SAP* (2014). [2](#)
- [CHA*14] CHAPIRO A., HEINZLE S., AYDIN T., POULAKOS S., ZWICKER M., SMOLIC A., GROSS M.: Optimizing stereo-to-multiview conversion for autostereoscopic displays. *Comp. Graph. Forum* 33, 2 (2014). [1](#), [2](#)
- [DKR*14] DABALA L., KELLNHOFFER P., RITSCHER T., DIDYK P., TEMPLIN K., MYSZKOWSKI K., ROKITA P., SEIDEL H.-P.: Manipulating refractive and reflective binocular disparity. *Comp. Graph. Forum* 33, 2 (2014), 53–62. [2](#)
- [DMHG13] DU S.-P., MASIA B., HU S.-M., GUTIERREZ D.: A metric of visual comfort for stereoscopic motion. *ACM Trans. Graph.* 32, 6 (2013). [2](#)
- [DRE*11] DIDYK P., RITSCHER T., EISEMANN E., MYSZKOWSKI K., SEIDEL H.-P.: A perceptual model for disparity. *ACM Trans. Graph.* 30, 4 (2011). [2](#)
- [DRE*12a] DIDYK P., RITSCHER T., EISEMANN E., MYSZKOWSKI K., SEIDEL H.-P.: Apparent stereo: The consweat illusion can enhance perceived depth. In *Proc. IS&T SPIE HVEI* (2012), pp. 1–12. [2](#)
- [DRE*12b] DIDYK P., RITSCHER T., EISEMANN E., MYSZKOWSKI K., SEIDEL H.-P., MATUSIK W.: A luminance-contrast-aware disparity model and applications. *ACM Trans. Graph.* 31, 6 (2012). [1](#), [2](#)
- [IKA07] ICHIHARA S., KITAGAWA N., AKUTSU H.: Contrast and depth perception: Effects of texture contrast and area contrast. *Perception* 36, 5 (2007), 686. [2](#)
- [KM14] KERSTEN D., MAMASSIAN P.: Cast shadow illusions. *Oxford Compendium Of Visual Illusions* (2014). (to appear). [2](#)
- [KZP*13] KIM C., ZIMMER H., PRITCH Y., SORKINE-HORNUNG A., GROSS M. H.: Scene reconstruction from high spatio-angular resolution light fields. *ACM Trans. Graph.* 32, 4 (2013), 73. [5](#)
- [LB99] LANGER M. S., BÜLTHOFF H. H.: Depth discrimination from shading under diffuse lighting. *Perception* 29, 6 (1999), 649–660. [2](#)
- [LHW*10] LANG M., HORNUNG A., WANG O., POULAKOS S., SMOLIC A., GROSS M.: Nonlinear disparity mapping for stereoscopic 3d. *ACM Trans. Graph.* 29, 3 (2010). [2](#)
- [MP89] MEDINA PUERTA A.: The power of shadows: shadow stereopsis. *JOSA A* 6, 2 (1989), 309–311. [1](#), [2](#)
- [MWA*13] MASIA B., WETZSTEIN G., ALIAGA C., RASKAR R., GUTIERREZ D.: Display adaptive 3d content remapping. *Comp. & Graph.* 37, 8 (2013), 983 – 996. [2](#)
- [RBD06] RUSINKIEWICZ S., BURNS M., DECARLO D.: Exaggerated shading for depicting shape and detail. In *ACM Trans. Graph.* (2006), vol. 25, ACM, pp. 1199–1205. [2](#)




Figure 7: Each set shows, from left to right: (i) no disparity and no shading stereo, (ii) some shading stereo, but no disparity, (iii) some disparity, but no shading stereo, and (iv) the same disparity, augmented by shading stereo. The insets represent the differences between the left and right eye images for each image.




Figure 9: Objects with sharp piecewise-planar features can sometimes generate color contrasts through color variation quickly, leading to a disturbing viewing experience. Additionally, planar objects will sometimes not benefit from shading stereo, as self-shading might not change significantly with light disparity (L) Faulty geometry (M) for live-action scenes can lead to misplaced shading, degrading image quality (R)

- [RFBW07] RAMANARAYANAN G., FERWERDA J., WALTER B., BALA K.: Visual equivalence: towards a new standard for image fidelity. *ACM Trans. Graph.* 26, 3 (2007). 4
- [SKHB11] SHIBATA T., KIM J., HOFFMAN D. M., BANKS M. S.: The zone of comfort: Predicting visual discomfort with stereo displays. *J. of Vision* 11, 8 (2011), 1–29. 1, 2
- [SLV*13] SCHER S., LIU J., VAISH R., GUNAWARDANE P., DAVIS J.: 3d+2dtv: 3d displays with no ghosting for viewers without glasses. *ACM Trans. on Graph.* 32, 3 (2013). 2
- [ŠPV11] ŠOLTÉSZOVÁ V., PATEL D., VIOLA I.: Chromatic shadows for improved perception. In *Proc. ACM/EG NPAR* (2011), pp. 105–116. 2
- [TDR*12] TEMPLIN K., DIDYK P., RITSCHER T., MYSZKOWSKI K., SEIDEL H.-P.: Highlight microdisparity for improved gloss depiction. *ACM Trans. Graph.* 31, 4 (2012), 92. 2, 3
- [vdPJD*14] VON DER PAHLEN J., JIMENEZ J., DANVOYE E., DEBEVEC P., FYFFE G., ALEXANDER O.: Digital Ira and beyond: Creating real-time photoreal digital actors. In *ACM SIGGRAPH 2014 Courses* (2014), SIGGRAPH '14, p. 1:384. 3
- [VMB*14] VANGORP P., MANTIUK R. K., BAZYLUK B., MYSZKOWSKI K., MANTIUK R., WATT S. J., SEIDEL H.-P.: Depth from hdr: Depth induction or increased realism? In *Proceedings of the ACM Symposium on Applied Perception* (2014), SAP '14, pp. 71–78. 2
- [YZWH12] YANG X. S., ZHANG L., WONG T.-T., HENG P.-A.: Binocular tone mapping. *ACM Trans. Graph.* 31, 4 (2012), 93. 2