

Highlight Microdisparity for Improved Gloss Depiction

Krzysztof Templin Piotr Didyk Tobias Ritschel Karol Myszkowski Hans-Peter Seidel
MPI Informatik



Figure 1: Physically-based highlights often cause binocular conflicts (left). On-surface highlights look less glossy and less authentic (middle). When rendered using the proposed technique, they are detached from the surface but do not introduce conflicts (right). Please note, that all anaglyph stereo images in the paper serve only as a preview of the effect. Refer to additional materials for high-quality stereo pairs.

Abstract

Human stereo perception of glossy materials is substantially different from the perception of diffuse surfaces: A single point on a diffuse object appears the same for both eyes, whereas it appears different to both eyes on a specular object. As highlights are blurry reflections of light sources they have depth themselves, which is different from the depth of the reflecting surface. We call this difference in depth impression the “highlight disparity”. Due to artistic motivation, for technical reasons, or because of incomplete data, highlights often have to be depicted on-surface, without any disparity. However, it has been shown that a lack of disparity decreases the perceived glossiness and authenticity of a material. To remedy this contradiction, our work introduces a technique for depiction of glossy materials, which improves over simple on-surface highlights, and avoids the problems of physical highlights. Our technique is computationally simple, can be easily integrated in an existing (GPU) shading system, and allows for local and interactive artistic control.

Links: [DL](#) [PDF](#) [WEB](#) [VIDEO](#) [DATA](#)

1 Introduction

Stereoscopic display technology – after waves of popularity in the 1950s and 1980s – has recently gained renewed attention. It is present in feature film, has been adapted by the gaming industry and found its way into scientific visualization. However, the possible range and variation of depth in the image are limited by viewing comfort considerations. In practice, a tradeoff between comfort and depth impression is achieved by disparity manipulation, e. g., compression. Such manipulations assume that disparity is well-defined by the scene’s geometry. While this assumption is valid

for solid diffuse surfaces, it no longer holds for view-dependent shading. Highlights, which are blurry images of the light sources, have their own depth, different from the surface on which they are visible (Fig. 2), and they are potential source of excessive disparities. Additionally, depending on the geometry of the reflecting surface, highlights can change shape, disappear, produce vertical disparities or result in different specular flow in both eyes. One solution to this problem is to assume a common (cyclopean) eye position when shading the surface. Doing so, we avoid conflicts, however highlights seem to be painted onto the surface. This is a significant shortcoming, as it is known, that highlight disparity is a strong factor for the perception of gloss [Blake and Bülthoff 1990; Hurlbert et al. 1991; Sakano and Ando 2010] and material authenticity [Wendt et al. 2008]. Nevertheless on-surface highlights are quite common, presumably for three main reasons: Because artists consider them to be more pleasant to watch; because of performance (e. g., in games that can not afford to shade twice [Sousa et al. 2012]); and because the necessary information is missing (e. g., in 2D-to-3D conversion).



Figure 2: An observer (a) looks at a specular surface (b) and sees a highlight (c) which is a (possibly blurred) image of a light source (d). Depending on geometry, the highlight appears behind (left) or in front of the surface (right). Note, that it is not placed on the surface.

The contributions of our paper are as follows:

- A problem analysis of highlight stereo depiction.
- A simple and safe alternative approach to highlight rendering that improves over on-surface and physical highlights.
- A perceptual study validating our approach.

2 Previous Work

Kirschmann [1895] was first to note the relation of highlights and disparity. Dove [1851] described the effect of highlight disparity on material perception and termed the resulting appearance “luster”. A key component for the distinct appearance of highlights seems to be binocular rivalry: Even when accounting for binocular disparity

© ACM, (2012). This is the author’s version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Graphics (Proc. SIGGRAPH) 31, 4, (2012). <http://doi.acm.org/10.1145/2185520.2185588>

and registering the two images, the resulting luminance does not agree in both eye’s images [Brewster 1861; Paille et al. 2001].

Blake [1985] derives equations that allow a machine – but maybe also the HVS – to infer shape from specular highlights. Later, the perception of highlight disparity was analyzed by Blake and Bülthoff [1990] in a matching experiment where participants were asked to adjust a rendered highlight’s disparity on a convex surface to obtain maximal realism. They found, that highlights were mostly placed behind the surface, i. e., physically, but with a bias towards the surface. We hypothesize that this bias originates from the fact, that looking at highlights (also real ones) with significant disparity causes visual discomfort. One could draw the conclusion, that whereas on-surface highlights are not realistic and reduce perceived glossiness [Wendt et al. 2008], rendering them physically is also not the best option regarding comfort. Hurlbert et al. [1991] observed that for convex surfaces glossiness perception is not sensitive to the relative highlight disparity with respect to the surface disparity, but highlights must be placed behind the surface. For concave surfaces perceived glossiness increases with the relative disparity irrespectively of its sign, which means that highlight placement behind the surface is acceptable, although usually it is not physically correct [Blake and Brellstaff 1988]. In the matching experiment conducted by Blake and Bülthoff [1990] the participants claimed the most realistic gloss impression is obtained by placing highlights on the concave surface or behind it. Our highlight microdisparity technique relies on these findings as we tend to place highlights behind the surface, while avoiding excessively large disparities so that best viewing comfort and realistic appearance is achieved.

Obein et al. [2004] investigated the relation between gloss sensation and specular gloss value in the context of monocular and binocular vision. They reported that binocular factors play the most important role in the judgment of high-gloss values, while for medium-gloss surfaces the gloss sensitivity is similar as for monocular vision. Clearly, for higher gloss values the distinctness of reflected image with high spatial frequency content enables better localization of relative highlight positions for each eye, which facilitates stereo matching [Hess et al. 1999]. Our highlight microdisparity technique conforms to these observations and focuses on highly glossy surfaces, where low relative disparities lead to realistic and comfortable gloss depiction.

In computer graphics, attempts have been made to achieve a perceptual normalization of gloss in monocular images [Pellacini et al. 2000; Wills et al. 2009]. On current stereo equipment, disparity is more limited than in real world scenes, and thus it requires special processing [Lang et al. 2010; Didyk et al. 2011]. However, existing work deals with diffuse surface disparity, and ignores the disparity of reflections. Some form of highlight processing is probably used in film production [Robertson 2009], however we are not aware of any technical publication analyzing the problem or suggesting a solution.

3 Our Approach: Highlight Microdisparity

Next, we will describe our approach to stereo highlight depiction. We will assume a shading model composed of a view-independent diffuse and a view-dependent specular component (e. g., Phong). First, the diffuse component is rendered with usual disparity into a stereo image pair. Second, the specular component is rendered into another stereo image pair using a single (cyclopean) eye position for shading. Consequently, highlights appear *on* the surface in this specular stereo image pair, because the specularities are calculated in a view-independent way, and thus they have zero-disparity relative to the surface. Finally, the specular stereo image pair is warped horizontally (i. e., its left sub-image to the left; the right sub-image

to the right). This warping re-introduces disparity between diffuse and specular shading. Horizontal warping avoids unpleasant vertical disparities found in physical stereo highlights [Blake 1985]. The warping is performed independently in each image of the specular stereo image pair. We describe how to control the per-pixel warping using a combination of basic microdisparity w_b , curvature w_c , edge detector w_e and artistic control w_a in the next four paragraphs.

Basic warping The warping constant w_b is chosen to be large enough to make highlights visibly detached from the surface but small enough not to introduce objectionable artifacts. Because the required amount of shift is small, the mismatch between geometry and highlights should not become apparent in monocular images. Our results were rendered using 2×2 super-sampling, and the warping was performed before scaling down the image. In every example in this paper, the shift of the highlights after down-sampling is not greater than 2 pixels for each view (4 pixels in total).

Curvature, edges and creases Highlight disparity depends on surface curvature: For high curvature, highlight disparity decreases because the surface “sees” larger portions of environment and the reflections undergo compression [Blake 1985; Fleming et al. 2004]. Choosing a constant disparity without accounting for curvature would lead to unrealistic results where highlights seem detached, when they should be placed almost on the surface. The purpose of the curvature factor w_c is to suppress the warping for highly curved surfaces, and it is proportional to the magnitude of the second derivative of the surface depth in horizontal direction, calculated in image space. This approach is inspired by the method of Vergne et al. [2009]. First, for every pixel \mathbf{p} , we determine the first derivative in horizontal direction using the normal vector $n(\mathbf{p})$: $g_x(\mathbf{p}) = -n_x(\mathbf{p})/n_z(\mathbf{p})$, where the z -axis points away from the camera. Next, we approximate the second derivative by $h_x(\mathbf{p}) = (g_x(\mathbf{p}_+) - g_x(\mathbf{p}_-))/2$, where \mathbf{p}_+ and \mathbf{p}_- are the horizontal neighbors of \mathbf{p} . Finally, we set:

$$w_c(\mathbf{p}) = \begin{cases} 1 & \text{if } |h_x(\mathbf{p})| = 0, \\ 0 & \text{if } |h_x(\mathbf{p})| \geq c_{\max}, \\ (c_{\max} - |h_x(\mathbf{p})|)/c_{\max} & \text{otherwise.} \end{cases}$$

In our experiments $c_{\max} = 0.03$ was used.

Another factor limiting the warping procedure are edges and creases, because the highlights should not move over them. We detect edges by convolving the image depthmap with a 3×3 Laplacian kernel, and thresholding the outcome. Thus, w_e equals 0 when an edge has been detected and 1 otherwise. The detection of creases is handled implicitly by the curvature weighting mechanism, since the second derivative has high magnitude around them.

Artistic control Spatially localized artistic control can be included by defining m sparse specular disparity constraints $(h_1, \epsilon_1), \dots, (h_m, \epsilon_m)$ at surface locations $\mathbf{p}_1, \dots, \mathbf{p}_m$. Gaussian radial basis functions are used to propagate the constraints to arbitrary spatial locations \mathbf{p} :

$$s(\mathbf{p}) = \sum_{i=1}^m e^{-\epsilon_i r_i^2} h_i \text{ with } r_i = |\mathbf{p} - \mathbf{p}_i|.$$

The parameters ϵ_i control the range of the constrains, whereas h_i – their strength and direction (an increase of the highlight disparity for positive h_i , and a decrease for negative h_i). The s function is evaluated for every pixel using a fragment shader. Finally, we set $w_a(\mathbf{p}) = 2^s(\mathbf{p})$, to approximately linearize the strength of the effect. Fig. 7 illustrates how the four weights influence the result. An example of manual changes to highlights is also given in Fig. 4. In all remaining pictures we assume $w_a = 1$.

Warping by gathering Having computed the required factors we combine them into a single warping coefficient: $w(\mathbf{p}) \in \mathbb{R}^2 \rightarrow \mathbb{R} = w_b w_c(\mathbf{p}) w_e(\mathbf{p}) w_a(\mathbf{p})$, that defines the warping map as

$$\bar{w}(x, y) = \pm \min_{i \geq 0} \{w(x \pm i, y) + i\}$$

for left and right views respectively. The map $\bar{w}(\mathbf{p})$ can be computed from w by checking a few pixels in the neighborhood of \mathbf{p} . Finally, the warped specular image is defined as:

$$I_S^W(x, y) = I_S(x + \bar{w}(x, y), y).$$

4 Results

We implemented our approach using a GPU to compute the warping and apply it to interactive rendering on a consumer PC. We report results for three use cases (full rendering, performance-critical rendering and 2D-to-3D) as well as a perceptual study.

4.1 Use cases

Full rendering Here, the full scene information is available, and the resources are sufficient to compute physically-based highlights, however our method is used in order to minimize unpleasant effects of physical reflections such as excessive horizontal disparities, vertical disparities or other binocular conflicts. The results of our approach in full rendering are shown in Fig. 3. The ability to locally and interactively control our approach is demonstrated in Fig. 4 and the accompanying video.

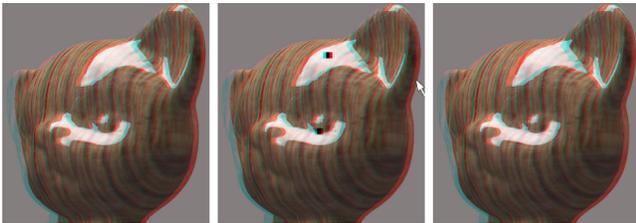


Figure 4: Highlight disparity (left) is locally adjusted near the ear and the eye (center) to obtain an improved result (right).

Performance-critical rendering Here, the full scene information is available as well, but the computational resources are too limited to produce two images fast enough e. g., for a computer game [Sousa et al. 2012]. Instead, the stereo image is produced using image-based warping techniques. Typically, the highlights would be warped together with the geometry, and would appear on-surface. Our technique can be used, to warp the highlights independently. An example usage in a game-like environment is shown in Fig. 5.

2D-to-3D In this use-case, our technique is an additional step in a 2D-to-3D pipeline which increases realism of the results obtained. First, the depth information in the picture is recovered, and the highlights are separated. Next, the diffuse and specular layer of the image are warped according to the depth buffer to produce a stereo pair. Then, small disparity is added to the specular layer, and both layers are combined. The result of this approach is presented in Fig. 6. The diffuse and specular layers were taken from [Tan and Ikeuchi 2005], the depth buffer was painted manually and constant normal field was assumed.

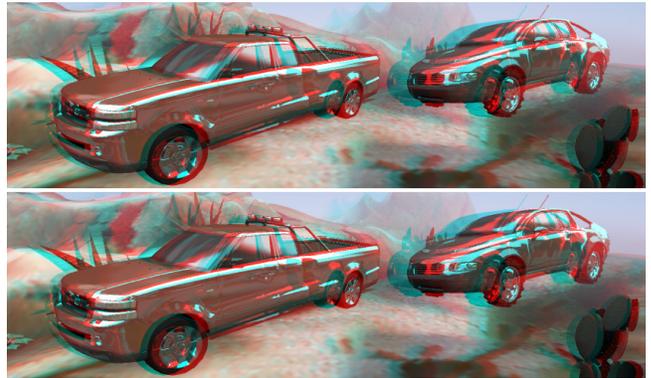


Figure 5: Performance-critical applications like games produce stereo images using image-based warping (top). Our approach can improve highlight depiction by warping them differently (bottom).

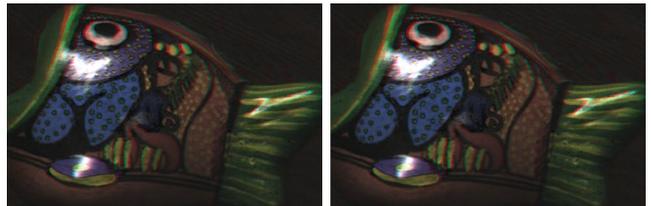


Figure 6: A stereo pair is generated from a single image by warping it according to depth (left). In addition to the basic warp, a small disparity is added to the highlights to enhance the picture (right).

4.2 Perceptual study

To verify our findings we conducted a perceptual study, where Figs. 1 and 3 were presented to 10 naïve subjects (7 F, 3 M) using a Zalman M240W polarized display. Three images of each scene with physical, on-surface and our highlights were shown on a neutral grey background next to each other. The placement of the versions were randomized for each test. In three sessions we asked subjects to indicate the most *unrealistic* highlight depiction, then the most *uncomfortable* impression, and finally the *preferred* one. We chose this negative formulation, as our main goal is to reduce artifacts while retaining material and stereo perception.

The results of the study are shown in Tab. 1. Surprisingly, many subjects judged physical simulation as unrealistic. However, as noted by Wendt [2008], a physically more correct rendering does not have to *appear* more realistic. While 38 % of the participants found on-surface highlights unrealistic, more subjects consider this technique superior to the costly, physical highlights, explaining its success in practical applications. In the last session 40 % were in favor of our technique, 34 % and 26 % preferred on-surface and physical highlights. We obtained $\kappa < 0$ (Fleiss' kappa) in all three sessions which suggests poor agreement between the subjects. The advantage of our technique in the first two sessions is statistically significant, however it is not significant in the last session.

Both negative formulation of the first two questions and the order of sessions may have biased the third experiment to our advantage. However, the construction of the study made our non-expert subjects more aware of possible issues with stereo gloss depiction. Surprisingly, in the preference session we did not find a significant effect. One may notice that our method works better when highlights are well-defined and isolated (Figs. 1, 3a–b), rather than complex or of lower sharpness (Fig. 3c–d). In the latter cases the conflict between left and right views is smaller, and the resulting discomfort is perhaps more tolerable.

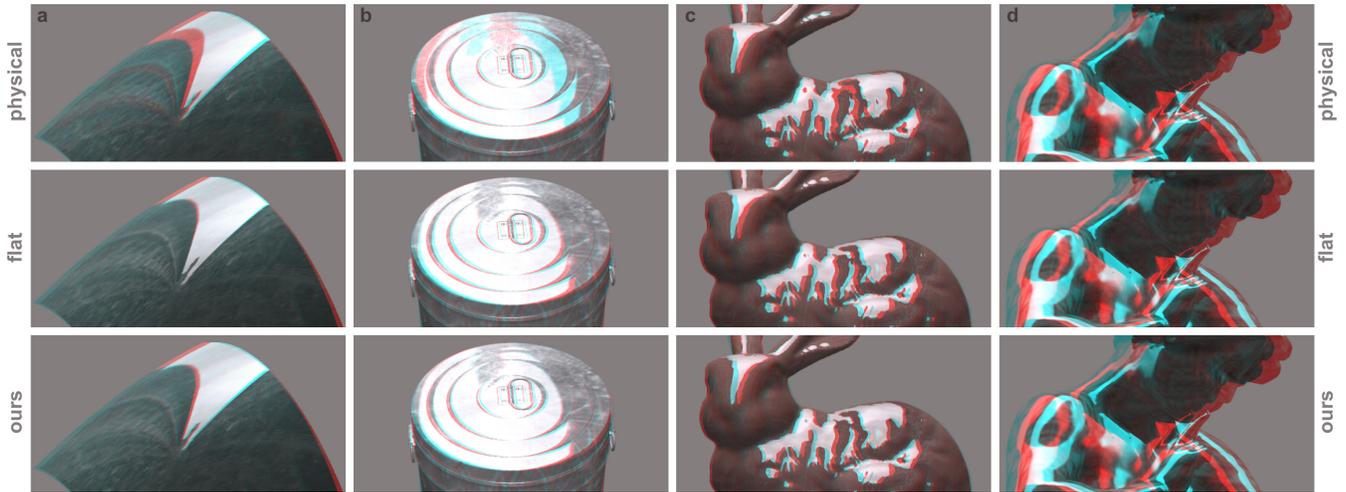


Figure 3: Results of using our method for different scenes. In each column, from top to bottom: physical highlights, on-surface highlights, highlights with microdisparity. Note that in print the images may be too small to see the effect. For best experience refer to the digital versions.

Fig.	Unrealistic			Discomfort			Preference		
	Phys.	Flat	Ours	Phys.	Flat	Ours	Phys.	Flat	Ours
1	60%	30%	10%	90%	0%	10%	10%	40%	50%
3 a	50%	20%	30%	70%	20%	10%	10%	30%	60%
3 b	70%	30%	0%	60%	30%	10%	30%	20%	50%
3 c	40%	50%	10%	60%	30%	10%	50%	30%	20%
3 d	30%	60%	10%	60%	40%	0%	30%	50%	20%
Avg.	50%	38%	12%	68%	24%	8%	26%	34%	40%
$p <$.0005	.0074		.0001	.0385		.1482	.3715	

Table 1: The results of the user study (please see the text).

5 Discussion

Including curvature when using our model proved to be useful. On an object with varying curvature (Fig. 7a), it is visible how constant highlight disparity would be objectionable. However, when the highlights appear mostly in high-curvature regions, the need for our technique is perhaps less obvious, because the conflicting highlights are of smaller size, and the disparity reintroduced by our method is attenuated by the curvature term (see Fig. 3c).

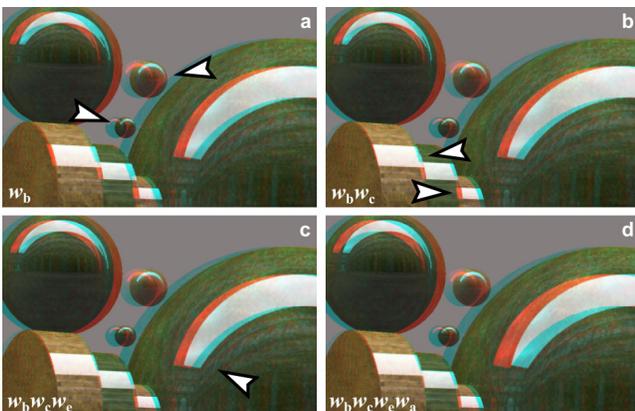


Figure 7: Without w_c , highlights on smaller balls are too detached (a). Missing w_e leads to jaggy and/or floating highlights on some edges (b). An artist may decide to adjust highlight disparity locally (c) to produce the optimal result (d).

We ignore the dependency of disparity visibility on surface glossiness (sharpness of highlights). For less glossy surfaces the amount of disparity introduced to highlights may be too small to be detectable, and thus the advantage over on-surface highlights negligible. On the other hand, for very blurry highlights even physical disparity can be hard to spot, and the overall difference between the three methods is not big in such cases. Fig. 3d is a good illustration to this issue: The highlights are blurry, so the disparity added by our method is not easily detectable. However, it is also hard to find a viewpoint in which physically-based highlights cause strong conflicts.

Another simplification is the lack of distinction between highlights and mirror reflections. In the case of reflections, the observer can distinguish shapes of the reflected objects, however disparity added by our algorithm lacks appropriate variation, and the cardboard effect appears (see Fig. 8). Those cases would need to be handled separately, either manually or by specialized algorithms. Finally, it

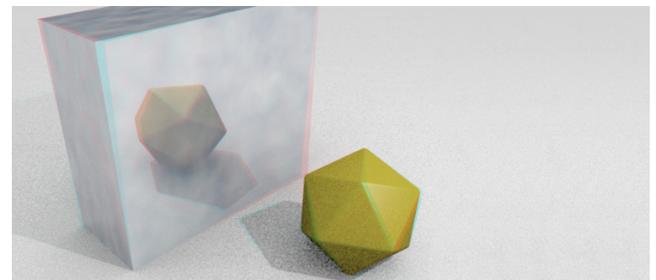


Figure 8: Failure case: In a planar mirror the reflected scene is clearly visible, but the reflection disparity is constant. Note the unnatural flatness of the polyhedron reflection (cardboard effect).

is possible that in many cases dealing with less pronounced highlight disparities than in our tests, physical computation will be superior to our method. Nevertheless, our method can be considered a safe automatic replacement for physical computation.

Temporal aspects Sakano and Ando [2010] investigated stereopsis with simultaneous head motion and observed only weak glossiness enhancement with respect to the binocular cue only. However, their stimuli did not contain any highlight disparity and were composed of flat, differently oriented facets. Wendt et al. [2010] considered smoothly curved surfaces and physically-correct highlight

disparity, which in combination with surface motion resulted in more reliable judgment of gloss strength than stereopsis for static stimuli. We relegate as future work more systematic study on the impact of camera and object motion, or dynamic lighting changes on the performance of our technique. As can be seen in the accompanying video, our method has very good temporal consistency.

Microrivalry The proposed method of introducing small highlight disparities can be seen as an instance of a more general class of techniques which we call *microrivalry*. There is a number of phenomena that lead to subtle differences in binocular images, such as thin-film interference, metal-flake paints or glare effects [Đuriković and Martens 2003; Ritschel et al. 2009]. Taking them into account during rendering stereoscopic images can increase their realism (Fig. 9).

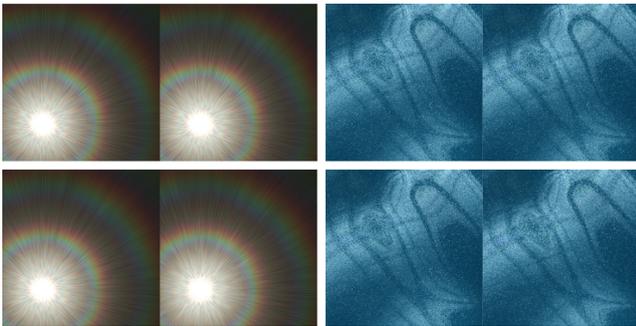


Figure 9: Images with subtle binocular conflicts (bottom) can look more realistic compared to ordinary images (top). The glare effect (left) appears independently in each eye. The sparkle patterns in metal-flake paints (right) are view-dependent.

6 Conclusion and Future Work

In this paper we presented a method of rendering highlights in stereoscopic 3D, that helps to preserve appearance of glossy materials. The technique is easy to include in an existing shading system and can be computed efficiently. Our approach provides a good alternative to physical and on-surface highlights.

Combining our method with depth-of-field, motion blur, and transparency remains future work. Other possibilities, besides accounting for the interplay with other depth cues, are changing the geometry to achieve a certain material appearance, eliciting a depth impression using only specular disparity [Howard 1995] or stylizing materials by exaggerating highlight disparity [Hurlbert et al. 1991].

References

BLAKE, A., AND BRELSTAFF, G. 1988. Geometry from specularities. In *Proc. Int. Conf. on Computer Vision*, 394–403.

BLAKE, A., AND BÜLTHOFF, H. 1990. Does the brain know the physics of specular reflection? *Nature* 343, 6254, 165–168.

BLAKE, A. 1985. Specular stereo. In *Proc. Int. J. Conf. on Artificial Intell.*, 973–976.

BREWSTER, D. 1861. On binocular lustre. *Reports of British Association* 2, 29–31.

DIDYK, P., RITSCHER, T., EISEMANN, E., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2011. A perceptual model for disparity. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4.

DOVE, H. 1851. Über die Ursachen des Glanzes und der Irradiation, abgeleitet aus chromatischen Versuchen mit dem Stereoskop. *Annalen der Physik* 159, 5, 169–183.

FLEMING, R. W., TORRALBA, A., AND ADELSON, E. H. 2004. Specular reflections and the perception of shape. *J Vision* 4, 9.

HESS, R., KINGDOM, F., AND ZIEGLER, L. 1999. On the relationship between the spatial channels for luminance and disparity processing. *Vis. Res.* 39, 3, 559–568.

HOWARD, I. 1995. Depth from binocular rivalry without spatial disparity. *Perception* 24, 67–67.

HURLBERT, A., CUMMING, B., AND PARKER, A. 1991. Recognition and perceptual use of specular reflections. *Investigative Ophthalmology & Visual Science* 32, 105.

KIRSCHMANN, A. 1895. *Der Metallglanz und die Parallaxe des indirecten Sehens*. Verlag von Wilhelm Engelmann.

LANG, M., HORNING, A., WANG, O., POULAKOS, S., SMOLIC, A., AND GROSS, M. 2010. Nonlinear disparity mapping for stereoscopic 3d. *ACM Trans. Graph. (Proc. SIGGRAPH)* 29, 4, 75.

OBEIN, G., KNOBLAUCH, K., AND VIENOT, F. 2004. Difference scaling of gloss: nonlinearity, binocularity, and constancy. *J Vision* 4, 9.

PAILLE, D., MONOT, A., DUMONT-BECLE, P., AND KEMENY, A. 2001. Luminance binocular disparity for 3d surface simulation. In *Proc. SPIE*, vol. 4299, 622.

PELLACINI, F., FERWERDA, J., AND GREENBERG, D. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proc. SIGGRAPH*, 55–64.

RITSCHER, T., IHRKE, M., FRISVAD, J. R., COPPENS, J., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2009. Temporal Glare: Real-Time Dynamic Simulation of the Scattering in the Human Eye. *Comput. Graph. Forum (Proc. Eurographics)* 28, 2, 183–92.

ROBERTSON, B. 2009. Monsters of the deep. *Comput. Graph. World* 32, 3.

SAKANO, Y., AND ANDO, H. 2010. Effects of head motion and stereo viewing on perceived glossiness. *J Vision* 10, 9, 15.

SOUSA, T., KASYAN, N., AND SCHULZ, N. 2012. *GPU Pro 3*. CRC Press, ch. CryENGINE, 163.

TAN, R., AND IKEUCHI, K. 2005. Separating reflection components of textured surfaces using a single image. *PAMI* 27, 2, 178–93.

ĐURIKOVIĆ, R., AND MARTENS, W. 2003. Simulation of sparkling and depth effect in paints. In *Proc. SCCG*, 207–213.

VERGNE, R., PACANOWSKI, R., BARLA, P., GRANIER, X., AND SCHLICK, C. 2009. Light warping for enhanced surface depiction. *ACM Trans. Graph. (Proc. SIGGRAPH)* 28, 3, 25.

WENDT, G., FAUL, F., AND MAUSFELD, R. 2008. Highlight disparity contributes to the authenticity and strength of perceived glossiness. *J Vision* 8, 1.

WENDT, G., FAUL, F., EKROLL, V., AND MAUSFELD, R. 2010. Disparity, motion, and color information improve gloss constancy performance. *J Vision* 10, 9.

WILLS, J., AGARWAL, S., KRIEGMAN, D., AND BELONGIE, S. 2009. Toward a perceptual space for gloss. *ACM Trans. Graph.* 28, 4, 103.