Specular Streaks in Stereo

Michael S. Langer School of Computer Science McGill University Montreal, Canada

langer@cim.mcgill.ca

Abstract—When a compact light source illuminates a horizontal shiny ground plane at an oblique angle, the resulting highlight is vertically oriented and highly elongated. We refer to such highlights as specular streaks. Specular streaks occur commonly on wet roadways, especially at night, for example the reflections of street lamps or car headlights. Here we present a 3D model of specular streaks seen by a binocular observer. We show that specular streaks produce binocular disparities that are consistent with 3D near-vertical columns of light beneath the roadway. We validate this model with results on both synthetic and real images. Specular streaks are an important problem for vision-in-badweather, in particular, for autonomous driving systems or driver assistance systems that rely on stereo to estimate scene depth.

I. INTRODUCTION

When a light source illuminates a horizontal shiny ground surface such as a wet street, polished floor, or body of water, it produces a specular highlight that is vertically elongated. Examples of such vertically elongated highlights are shown in Figure 1. We refer to such highlights as *specular streaks*.



Fig. 1. Examples of specular streaks. The light sources produce vertically elongated highlights on the wet ground region.

Specular streaks have been depicted by many artists in landscape paintings, such as in Van Gogh's *Starry Night Over the Rhone* and Monet's *Impression Sunrise*. Specular streaks have also been studied in optics; for example, the Dutch physicist Minnaert in his classical study of light phenomena in natural scenes (1937) referred to them as *light columns*[1]. Specular streaks have been discussed in computer graphics as well. For example, it is well known that the Blinn-Phong model correctly produces vertically elongated highlights on the ground, whereas the Phong model does not [2], [3].

Curiously, specular streaks have never been addressed in the computer vision literature. This is surprising since specular streaks occur commonly in an increasingly important scenario for computer vision, namely wet roadways especially at night. Street lamps, traffic lights, the headlights of other vehicles, David Bourque School of Computer Science McGill University Montreal, Canada david.bourque@mail.mcgill.ca

and illuminated signs all produce specular streaks. Computer vision systems will need to deal with specular streaks in order to estimate depth robustly in these scenarios.

More generally, specular *highlights* have always been challenging for 3D reconstruction methods, especially using binocular stereo or structure from motion. The reason specular highlights are challenging is that do not produce reliable correspondences for triangulation and hence can easily produce spurious depth estimates. We would therefore expect specular streaks to pose a similar problems for these methods. An important application is driver assistance systems and vision for autonomous vehicles which may rely on binocular stereo (among other sensors such as lidar) to estimate depth and scene layout. We argue that for such systems to be effective, they need to take account of the peculiar and particular phenemenon of specular streaks.

Our main objective in this paper is to develop a basic appearance model of specular streaks that are viewed by a binocular observer. Specular streaks are not merely a 2D image phenomenon. Rather they have 3d properties. Specular streaks have binocular disparities and hence can give rise to spurious depth estimates that are different from the ground plane. We model the 3D properties of specular streaks, in particular their binocular disparity and depth.

An overview of the paper is as follows. In Section II we review some background. We review the related computer vision problem of 'shape from specularity'; we also review the basic geometry of specular streaks, namely how they arise from microfacet models of glossy surfaces. The novel contributions begin in Section III. We present a 3D model of a specular streak from a single point light source, and we show that such a specular streak defines a virtual light contour below the ground plane and that this contour is near vertical. This is the 3D version of Minnaert's (2D) light column mentioned above. In Section IV, we use the model to predict the binocular disparity of points on a specular streak. We test this binocular disparity model using a stereo pair of rendered specular streaks. We show that the binocular disparities that are estimated from the specular streak in stereo is approximately constant along the streak. We also give examples of estimated binocular disparities from a real image including specular streaks and other vertical objects. We conclude in Sec. V.

II. PREVIOUS WORK

A. Specular streaks in single images

Specular streaks in single images can be modelled in several ways. To our knowledge, Minnaert [1] was the first to model specular streaks in single images. He considered a wet horizontal ground plane with small variations in surface normals, *e.g.* a water surface with small waves or a very wet street. He derived bounds on the shape of a specular streak in an image, namely on the visual angle defining the height and width of a streak in an image. These bounds depended on the maximum angular deviation (in radians) of the local surface normal on the ground surface, namely the angle away from the gravity direction.

An alternative and more common model of specular reflection that gives rise to specular streaks in single images uses *probabilities* of surfaces normals, where the surface is assumed to consist of *microfacet* mirrors whose normals have some random distribution [4], [5]. Such a microfacet model is the basis of the Blinn-Phong and other models in computer graphics [6]. The Blinn-Phong model is often written as

$$I_{Blinn} = k_a + k_d \max(0, \mathbf{n} \cdot \mathbf{l}) + k_s (\mathbf{n} \cdot \mathbf{h})^s \tag{1}$$

where the three terms are ambient, diffuse, and specular, and the latter produces the specular highlight. Here l is the unit length vector in the direction of a point light source which in the above equation is assumed to be at infinity; **n** is the macroscopic surface normal which in the case of specular streaks is the constant ground plane normal; **h** is the "halfangle" vector:

$$\mathbf{h} = \frac{\mathbf{v} + \mathbf{l}}{\|\mathbf{v} + \mathbf{l}\|} \tag{2}$$

between the unit light **l** and unit viewer **v** directions. The half angle vector **h** is the surface normal of the microfacet that would produce a mirror reflection from the light source to the viewer. (We will use the half angle vector later in our model – see Eq. 4.) The exponent s in Eq. (1) is the *shininess*. It captures the surface roughness, or spread of the distribution of microfacet normals.

In more general model, the microfacet distribution appears as a term within a bidirectional reflectance distribution function (BRDF) $\rho(\omega_i, \omega_r)$ [7]. The reflected radiance $L_r(\omega_r)$ leaving an isotropic surface in direction ω_r and from a point **X** can be written

$$L_r(\mathbf{X}, \omega_r) = \int \rho(\mathbf{x}, \omega_i, \omega_r) \ L_i(\omega_i)(\mathbf{n} \cdot \omega_i) \ d\omega_i \qquad (3)$$

where $L_i(\omega_i)$ is the incident radiance arriving from direction ω_i , and ω_r is a reflection direction toward the viewpoint. In this general model, the light source does not need to be at infinity, nor does it need to be modelled as a point. An example is the Ward model [8]. For the isotropic reflectance case, the Ward model can be written as the sum of a matte and glossy (specular) term:

$$\rho_{Ward}(\omega_i, \omega_r) = \frac{k_d}{\pi} + \frac{k_s}{\sqrt{\cos \omega_i \cos \omega_r}} \frac{e^{-\tan(\delta)^2/\alpha^2}}{4\pi\alpha^2}$$

where k_d diffuse reflectance, k_s is the overall specular reflectance, and the microfacet distribution is captured by the Gaussian term where $\delta = acos(\mathbf{n} \cdot \mathbf{h})$ is the angle between normal \mathbf{n} and \mathbf{h} . The constant α controls the spread of the microfacet normals, and hence is related to the surface roughness. Many other BRDF models for glossy surfaces have been proposed such as the Cook and Torrance model [9] and GGX model [10]. These models attempt to capture various geometric and physical effects more explicitly and precisely, namely the self-shadowing and occlusions by the microfacts and Fresnel effects [2].

Figure 2 shows three renderings of scenes containing three light sources each. These light sources are at depth Z_L from the viewer (units arbitrary). Figure 2 (a) and (b) use the Blinn-Phong and Ward models respectively with a point source approximation for each source. Figure 2 (c) uses a ray sampling model with non-point spherical source. The surface roughness parameters were chosen so that the three models produce qualitatively similar specular streak effects, namely vertically elongated highlights of similar width and length. For the ray sampled model, a finite light source (rather than point light source) is used, and so the specular streaks are slightly wider. See Figure 2 caption for more details.



Fig. 2. Examples of renderings using three models: (a) Blinn-Phong, (b) Ward, and (c) ray sampled. The horizontal field of view angle is 14 deg and the vertical angle is 28 deg. The thin horizontal gray line in the middle of each image indicates the horizon. The three light sources are at heights $\frac{Z_L}{50}, \frac{3Z_L}{50}, \frac{3Z_L}{50}$ where Z_L is the depth of these sources. The viewer is at the same height $\frac{2Z_L}{50}$ as the middle source, and so the middle source in each plot appears on the horizon. For (a), the shinyness *s* was 1000 and for (b) $\alpha = .05$. For (a) and (b) a point source model is used. For (c), the light source has finite radius $\frac{Z_L}{250}$. To more easily compare the shape of the streaks, we scaled the intensity of the three lights (proportionally to light height) when rendering each streak, which gave the three streaks roughly the same maximum intensity. The rendered image intensity of the lights themselves was clipped to the maximum intensity value along the streaks. All three models produce roughly the same qualitative effect.

B. Shape from Specularity

Despite their common presence in natural roadway scenes in bad weather, specular streaks have not been discussed in computer vision. Rather, most work on specular reflection in computer vision has addressed the case of shiny surfaces that are smooth and compact (objects). Here we briefly review some of this work. The classical 'shape from specularity' problem is to estimate the 3D surface geometry of a shiny or mirror surface. Since our work is concerned with the binocular appearance of specular streaks, we begin our review with shape from specularity studies that addressed the stereo vision problem. Blake [11] was the first to formulate this problem, observing that a specular highlight on a convex or concave surface defines a virtual light source behind or in front of the surface, respectively, and so the disparity of the highlight relative to the surface itself could provide a cue to the sign of surface curvature [12]. A similar idea was applied to a moving observer [13], and it was shown that quantitative shape recovery was possible under some conditions. For example, if an observer knows its own motion and the position of the light source, then the observer can recover the 3D space curve travelled by the specular highlight along the surface [13], [14].

The above arguments do not require that the specular reflections are highlights, namely reflection of light sources. Rather they require only a specular reflection of some pattern that can be tracked across multiple image frames. Subsequent works often considered curved mirrored surfaces, which yield a deformed mirror image of the surrounding environment [15], [16]. For the stereo (two frame) case, one can ask under what conditions there exist matching points between the left and right views [17], [18], For example, when is an environment point visible in both the left and right mirror images, and when are such matches unique? Unlike in the Lambertian case, matching points do not necessarily lie on epipolar lines, and so there may not be well-defined 3D virtual point for the match. For most models of 3D shape estimation from specular motion or disparity, the observer needs to make strong assumptions either about the surface shape [19] or the viewing conditions, for example that the viewer and light source are distant relative to the size of the object[20]. Some methods typically assume the environment that is reflected in the mirror is a known calibration pattern, either at a fixed position [21] or unknown position[22].

The specular *streak* scenario that we introduce is quite different from the scenarios addressed above, which all concerned the case of a compact curved surface shape that is viewed from a nearby position under typically very controlled conditions. The viewing scenario for specular streaks is quite different. They arise when a glossy ground surface is illuminated at a high incident angle and is viewed from a large reflection angle. In the next section, we develop a model of a specular streak as a virtual 3D light column which is roughly vertical and located beneath the light source.

III. SPECULAR STREAKS AS 3D LIGHT COLUMNS

In this section we examine the basic appearance phenomenon, that specular streaks appear as 3D vertical light columns. Figure 3 shows a "texture map" representation of the specular streaks in Figure 2, namely the same image intensities in Figure 2 are shown as a function of ground plane coordinates (X, Z) rather than image coordinates (x, y). These

texture map images were computed using the homography [23]

$$(x,y) \leftrightarrow (\frac{fX}{Z}, \frac{fh}{Z})$$

between image and ground plane coordinate systems. Note the specular streaks in the texture map are much more elongated than the streaks in the image plane; specular streaks occur in images *despite* perspective foreshortening, not because of it.

Since the specular streaks in the Figure 2 image are all vertically oriented, they meet at a vanishing point at infinity, namely in the negative y direction in the image. Similarly in the texture map, the streaks are oriented such that their major axes appear to pass through the origin (X, Z) = (0, 0), which is the ground plane point at the viewer's feet. That point is the apex of the truncated pyramid(s) shown. These observations are all consistent with the specular streaks being interpreted as vertical columns in 3D. As we show next, specular streaks are indeed geometrically consistent with 3D light columns that are approximately vertical, located below the light source and beneath the ground plane.



Fig. 3. Texture maps defined on the (X, Z) plane, corresponding to the specular streak images of the ground plane regions in Fig. 2. The trapezoid shows the region of the ground plane that is within the observer's 14 deg field of view and is up to the depth Z_L of the light source. The four edges correspond to the boundary of the lower halves of the reflected light components of the images in Fig. 2. The apex of each (truncated) triangle is the viewer XZ position (0, 0).

Let the scene contain a ground plane Y = 0 and a point light source at position

$$\mathbf{X}_L = (X_L, Y_L, Z_L).$$

So Y_L is the height of the source above the ground plane. If the ground plane were a perfectly flat *mirror*, then there would be a single virtual point light source at position $(X_L, -Y_L, Z_L)$ namely the mirror image of the light source below the ground plane. Specular streaks arise when the ground plane is *glossy*. What can we say about the 3D virtual image of the source in this case? Let the viewer be located at Z = 0, and at the same XY position as the light source, so the viewer is at $\mathbf{X}_V = (X_V, Y_V, 0)$. To define the specular streak, assume a microfacet model of glossy reflectance. At each ground plane position $\mathbf{X} = (X, 0, Z)$, there is a single potential microfacet orientation \mathbf{h} – see Eq. 2 – such that a ray from the point

source is mirror reflected to the viewer. For this **X** and **h** pair, a virtual light position is defined by mirror-reflecting the light vector $\mathbf{L} \equiv \mathbf{X}_L - \mathbf{X}$ about the unique plane geometric that passes through **X** and has normal **h**. The position of the virtual light source is

$$\mathbf{L}_{virtual} = \mathbf{X} + (\mathbf{I} - 2\mathbf{h}\mathbf{h}^T)\mathbf{L}.$$
 (4)

So this virtual light source $\mathbf{L}_{virtual}$ depends both on \mathbf{X} and on the viewing position \mathbf{X}_V (via **h**).

Proposition: Suppose a point source is at X_L above a glossy ground plane, and the viewer moves along the line $X_V(s) = (X_V + s, Y_V, 0)$. Then for each depth Z between 0 and Z_L , there is a unique virtual light source $\mathbf{L}_{virtual}$ given by Eq. (4) that is invariant to the observer motion. (See the Appendix for the proof of this proposition.)

Figure 4 uses the model of Eq. (4) to compute the locus of $\mathbf{L}_{virtual}$ points for three different light sources (red, blue, and green) – see caption for details. Note that the virtual light columns are each slightly sloped, so they are not exactly vertical.



Fig. 4. (a) A viewer and three points sources are shown in a YZ plane (constant X). $\mathbf{L}_{virtual}$ point source positions are computed using Eq. (4). The viewer height is $\frac{Z_L}{10}$ and the three different light source heights are $\frac{Z_L}{20}, \frac{Z_L}{10}, \frac{3Z_L}{20}$. The plot shows the $X_V = X_L$ plane. We plot only those virtual light source positions whose corresponding **h** vector is less than 4 deg away from the gravity vector (Y axis).

Figure 5 illustrates why the virtual light columns are slightly sloped. We consider just one of the point sources (red). Three points along the ground plane are indicated along with the **h** vector of the microfacet that produces a mirror reflection of that point source to the viewer. When the half vector **h** is gravity aligned, the virtual source $\mathbf{L}_{virtual}$ in Eq. 4 is the mirror reflection of the point source below the ground. For depths Z that are closer to viewer, **h** is slanted slightly away from the viewer; this decreases the Z component of $\mathbf{L}_{virtual}$ and pushes it deeper below the ground plane. Similarly, for depths Z that are farther from the viewer, the microfacet **h** slants slightly toward the viewer, which increases the Z component of $\mathbf{L}_{virtual}$; this brings the virtual point closer to the ground plane.

A few caveats should be given about the above proposition and the figures. First, a microfacet with suitable half angle vector \mathbf{h} must be present for the streak to appear, and this is



Fig. 5. The slant of the virtual light column is due to the varying slopes of the microfacet half vectors \mathbf{h} that mirror reflect rays from the point light source to viewer.

certainly not guaranteed. Indeed the reason for the streak appearance is that, at most points on the ground, the probability of the necessary half-vector \mathbf{h} occuring is essentially 0.

Second, the proposition suggests that the locus of virtual light points $\mathbf{L}_{virtual}$ is a 1D curve in 3D space, namely one point per depth Z. However, the proof of the proposition (see Appendix) only considers those **h** vectors whose component h_X is 0. Other **h** vectors can contribute which is why specular streaks have a finite thickness even in the ideal case of a point light source, as Figure 2(a,b) showed.

Third, the proposition assumes a point light source. When the light source instead has a finite size, the specular streak will be thickened because it will be the superposition of specular streaks that are due to the component source points, as in the ray sampled model of Figure 2(c).

In the next section, we ask a related question: when we render a stereo pair of specular streaks, how well do the binocular disparities obey the model of Eq. 4? As we will see, the model is not perfect but it captures the main qualitative effect that the disparity (or depth) of the virtual 3D column is approximately constant. We also show some examples of specular streaks in real images, and we compute the binocular disparities of specular streaks and compare them with the disparities of other vertical objects.

IV. BINOCULAR DISPARITY ALONG A SPECULAR STREAK

From the model of Eq. 4, we expect that a specular streak defines a 3D virtual light column at approximately the same depth as the light source. Thus, we would expect that the binocular disparities along a specular streak are approximately constant, namely approximately equal to the disparity of the light source.

Let T_X be the interocular distance between two rectified cameras which are both pointing in the Z direction. Then the disparity in the image plane of corresponding image points for a 3D point at depth Z_l will be

$$disparity = \frac{fT_X}{Z_l} \tag{5}$$

where f is the distance to the projection plane.

Figure 6 shows the binocular disparities from a specular streak rendered in stereo using the computer graphics model (Ward), similar to the middle specular streak in Fig. 2. For

this example, the light source was at a Z distance of 15 m and height 1.5m above the ground, the viewer height was 1.5m, and the viewer's interocular distances was 6.5 cm. The values are correspond to a typical street scene that is seen by a binocular camera.

To estimate the binocular disparity from the stereo pair of images, we used a 1D Lucas-Kanade style least squares method. The results are show in the black curve. We plot one disparity value per image row, namely the estimated disparity at the center of the streak. We also plot the disparities computed using the theoretical model of Eq. 4, converting from depth to disparity using Eq. 5. See gray curve. The intersection point of the two curves corresponds to the point where there would be a virtual point light source (mirror reflection) in the case that the ground had a perfect mirror reflectance rather than glossy.

There are differences between the two curves, but these differences are very small in an absolute sense. (Note the small range of disparities on the abscissa.) For example, the disparity of the $\mathbf{L}_{virtual}$ locus (blue) varies over a range of .007 degrees over an elevation angle of nearly 9 degrees. The takeaway from this plot is that disparity is nearly constant (but not exactly constant) on both of these curves.



Fig. 6. Comparison of theoretical model of disparity from Eq. 4 along a specular streak (blue) and disparities computed (black) from a rendered stereo image of a specular streak. The intersection point corresponds to the position of the mirror reflection of the point source in the case that the ground surface had been a perfect mirror. *The disparity scale in this plot has been expanded to magnify differences between the curves – see text.*

Figure 7 show three examples of stereo images of rainy night street scenes. (The first two were shown in Fig. 1.) The images were shot with a Fujifilm FinePix Real 3D stereo camera. The images were then converted to grey level and histogram equalized to enhance the contrast of the nonspecular regions. The left image in each pair was shifted by hand so that points near the center of the scene had roughly 0 disparity. A very small vertical shear correction (the same for all three images) was also applied so that disparities were horizontal only.

Readers who can free-fuse the stereo pairs in the left two columns or who have access to red-cyan anaglyph glasses (and who can experience stereopsis) will be able to perceive the vertical columns of the specular streaks. For the analyph pairs, the 3D vertical columns of the streaks might be more salient if one inverts the images (and left-right swaps the red-cyan lenses!), since then the roadways become ceiling surfaces. Under normal viewing, our visual systems have a strong prior to see the ground surfaces as slanted upwards (floors) and for some viewers the streaks might be perceived as lying on the ground rather than beneath the ground, contrary to the actual disparity information.



Fig. 7. Stereo pairs from Fig. 1. The images have been converted to grey level and histogram equalized. The left two images can be free-fused; the right images requires red-cyan anaglyph glasses.

Figure 8 shows examples of stereo disparities computed from the images of Figure 7. For each example, regions are marked from the left image of each pairs. For each region and for each of the rows within the region, the horizontal disparity was estimated over a window centered on the central pixel that row, using a version of the Lucas-Kanade method.

In Figure 8 (top), four regions are selected. The disparities and hence depths are roughly constant within each region, and these constants different between regions. The 'poles' regions have constant depth because they are vertical, and the specular streak regions have constant depth as per our main argument. The ordering of these disparity constants are also roughly as expected. The disparity of the specular streak from the car headlight is positive since the car is in the foreground, whereas the disparity of the specular streak from the distant street lamp negative. (Recall that the images were shifted to put 0 disparity at roughly the center of the scene.) The near telephone pole has a constant disparity whose value lies between the disparities of the two specular streaks, since the pole's depth lies between the car headlamp and streetlamp. The region with two poles in the distance has the smallest disparity.

Figure 8 (middle) shows a second example, now with five selected regions. The largest disparity region is the foreground region containing part of a manhole cover. The disparity varies over that region, both because the texture and because the actual disparity varies along the ground. The disparity of the other regions is roughly as expected, decreasing as distance increases. There is an exception, however: we expected the



Fig. 8. Binocular disparity as a function of y position (vertical) for selected regions. See text.

disparity of the streaks to be closer to the disparity of the distant background pole. The reason that the disparities streaks may be greater than expected is that the wetness of the street varies spatially, and so the edge of the streak is partly defined by the wet-dry boundary of puddles. Such a material boundary is an image feature; it is located *on the ground* and so has the disparity of the ground. Such boundaries evidently complicate the appearance of specular streaks in general.

Figure 8 (bottom) shows a third example. Again, the ordering of the disparities corresponds to the ordering of depths in the scene. Note that the single specular streak in the red region is at the same y position in the image as the pole (magenta) but its disparity is quite different, since it is corresponds to a (virtual) vertical column below a streetlamp that is at the far end of the scene. Finally, the other specular streaks (green) are so far away that they have roughly the same disparity as all distant points, specular or not.

V. CONCLUSIONS

Specular streaks are an example of "vision in bad weather". Previous vision-in-bad-weather research has tended to address problems of dehazing, and rain and snow removal [24], [25]. Specular streaks are a naturally related problem, and they present similar challenges. In particular, specular streaks pose challenges for important application domains such as autonomous vehicles and driver assistance systems. Such vision systems should be robust to all problems of vision in bad weather, including specular streaks. Such vision systems should detect specular streaks, and either properly interpret them, remove [26] or ignore them. In particular, specular streaks obviously should not be interpreted literally as vertical columns below the ground, even though this is what their binocular disparities indicate.

As human observers, we typically do not perceive specular streaks as 3D vertical columns of light beneath the ground, unless we pay attention to them. This is similar to how we do not perceive shadows, unless we pay attention to them [27]. Our own vision systems seem to treat specular streaks for what they are, namely long glossy highlights, and we tend to ignore the geometric information (3D vertical columns) that they provide. However, for computer vision systems to ignore specular streaks in a similar way, they would need to be told to do so, or trained to do so. One way to train computer vision systems to do so would be to treat specular streaks as just another type of object in a roadway to detect, along with as pedestrians, cyclists, vehicles, sign posts, etc.. Vision systems would then by able to ignore specular streaks, rather than suffering any consequences from mistakes in interpretation of what the streak is, in particular, misinterpreting the streak as a gaping hole in the roadway ahead.

APPENDIX

Here we prove our claim in the Proposition that defines the 3D locus of a specular streak – see Eq. 4. Specifically we show that $\mathbf{L}_{virtual}$ is invariant to the X_V coordinate of the viewer. To show this, we first restrict ourselves to the case that the viewer and light source share the same X coordinate, $X_V = X_L$. Take the ground plane point $(X_L, 0, Z)$ at depth Z between the viewer and light source. There is a unique microfacet normal **h** at this point that could produce a mirror reflection from source to viewer. This **h** defines a virtual light source via Eq. 4. Moreover, this $\mathbf{L}_{virtual}$ and **h** would lie in this $X = X_L = X_V$ plane. We next show that this point $\mathbf{L}_{virtual}$ would also be a virtual light source point if the viewer were to move in the X direction.

Consider the plane Π that contains the above mirror reflection point $\mathbf{X} = (X_L, 0, Z)$ and has normal **h**. This plane Π intersects the ground plane along the *constant depth line* which can be written parametrically with parameter t as $(X_L + t, 0, Z)$. When the light rays from the point source reflect off points on this line, according to the microfacet normal **h** of plane Π , these reflected rays would (1) geometrically diverge from the virtual light source $\mathbf{L}_{virtual}$, and (2) these rays would all pass through the parametric line $(X_V+s, Y_V, 0)$ with parameter s, which contains the viewer position (s = 0). It follows that if the viewer were to move along this parametric line, then the above point $L_{virtual}$ would remain visible to the viewer, namely the viewer would intercept rays reflected off the microfacets with normal **h** on the line defined by the intersection of Π and the ground plane. This completes the proof.

REFERENCES

- M. Minnaert, *Light and Color in the Outdoors*, 5th ed. Springer-Verlag, 1993.
- [2] M. Pharr, W. Jakob, and G. Humphreys, *Physically Based Rendering: From Theory to Implementation*, 3rd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2016.
- [3] T. Akenine-Moeller, E. Haines, and N. Hoffman, *Real-Time Rendering*, *third edition*. A K Peters/CRC Press, 2008.
- [4] C. Cox and W. Munk, "Measurement of the roughness of the sea surface from photographs of the sun's glitter," *Journal of Optical Society of America*, vol. 44, p. 838–850, 1954.
- [5] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *Journal of the Optical Society of America*, vol. 57, no. 9, pp. 1105–1114, 1967.
- [6] J. F. Blinn, "Models of light reflection for computer synthesized pictures," in *SIGGRAPH*, 1977, pp. 192–198.
 [7] F. Nicodemus, "Directional reflectance and emissivity of an opaque
- [7] F. Nicodemus, "Directional reflectance and emissivity of an opaque surface," *Applied Optics*, vol. 4, no. 7, p. 767–775, 1965.
- [8] G. J. Ward, "Measuring and modeling anisotropic reflection," in SIG-GRAPH, 1992, pp. 265–272.
- [9] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," SIGGRAPH, vol. 15, no. 3, pp. 307–316, 1981.
- B. Walter, S. R. Marschner, H. Li, and K. E. Torrance, "Microfacet models for refraction through rough surfaces," in *Proceedings of the Eurographics Symposium on Rendering Techniques*, 2007, pp. 195–206.
- [11] A. Blake, "Geometry from specularities," in IJCAI, 1985.
- [12] A. Blake and H. H. Bulthoff, "Shape from specularities: computation and psychophysics," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 331, no. 237-252, Feb. 1991.
- [13] A. Zisserman, P. Giblin, and A. Blake, "The information available to a moving observer from specularities," *Image and Vision Computing*, vol. 7, no. 1, pp. 38–42, February 1989.
- [14] M. Oren and S. Nayar, "A Theory of Specular Surface Geometry," *International Journal on Computer Vision*, vol. 24, no. 2, pp. 105–124, 1996.
- [15] R. W. Fleming, A. Torralba, and E. H. Adelson, "Specular reflections and the perception of shape," *Journal of Vision*, vol. 2, no. 9, pp. 798– 820, 2004.
- [16] S. Savarese, M. Chen, and P. Perona, "Local shape from mirror reflections," *International Journal of Computer Vision*, vol. 64, no. 1, pp. 31–67, 2005.
- [17] A. A. Muryy, A. E. Welchman, A. Blake, and R. W. Fleming, "Specular reflections and the estimation of shape from binocular disparity," *Proceedings of the National Academy of Sciences*, vol. 110, no. 6, pp. 2413–2418, 2013.
- [18] A. A. Muryy, R. W. Fleming, and A. E. Welchman, "Key characteristics of specular stereo," *Journal of Vision*, vol. 14, no. 14, pp. 1–26, 2014.
- [19] S. Roth and M. Black, "Specular flow and the recovery of surface structure," in CVPR, vol. 2, 2006, pp. 1869–1876.
- [20] Y. Adato, Y. Vasilyev, T. Zickler, and O. Ben-Shahar, "Shape from specular flow," *IEEE Transactions Pattern Analysis and Machine Intelligence*, vol. 32, p. 2054–2070, 2010.
- [21] T. Bonfort, P. F. Sturm, and P. Gargallo, "General specular surface triangulation," in ACCV, 2006, pp. 872–881.
- [22] K. Han, M. Liu, D. Schnieders, and K. K. Wong, "Fixed viewpoint mirror surface reconstruction under an uncalibrated camera," *IEEE Trans. Image Process.*, vol. 30, pp. 2141–2154, 2021.
- [23] R. Hartley and A. Zisserman, Multiple View Geometry in Computer Vision. Cambridge University Press, 2003.
- [24] S. Nayar and S. Narasimhan, "Vision in bad weather," in *ICCV*, vol. 2, 1999, pp. 820–827 vol.2.
- [25] P. C. Barnum, S. Narasimhan, and T. Kanade, "Analysis of rain and snow in frequency space," *International journal of computer vision*, vol. 86, no. 2, pp. 256–274, 2010.
- [26] A. Artusi, F. Banterle, and D. Chetverikov, "A survey of specularity removal methods," in *Computer Graphics Forum*, vol. 30, no. 8. Wiley Online Library, 2011, pp. 2208–2230.
- [27] R. Casati and P. Cavanagh, The Visual World of Shadows. The MIT Press, 2019.