

Motion of specularities on low-relief surfaces: frequency domain analysis

Yousef Farasat and Michael S. Langer

School of Computer Science, McGill University
3450 University St. Montreal, Canada
Email: {yousef,langer}@cim.mcgill.ca

ABSTRACT

Typical studies of the visual motion of specularities have been concerned with how to discriminate the motion of specularities from the motion of surface markings, and how to estimate the underlying surface shape. Here we take a different approach and ask whether a field of specularities gives rise to motion parallax that is similar to that of the underlying surface. The idea is that the caustics that are defined by specularities exist both in front of and behind the underlying surface and hence define a range of depths relative to the observer. We asked whether this range of depths leads to motion parallax. Our experiments are based on image sequences generated using computer graphics and Phong shading. Using low relief undulating surfaces and assuming a laterally moving observer, we compare the specular and diffuse components of the resulting image sequences. In particular, we compare the image power spectra. We find that as long as the undulations are sufficiently large, the range of speeds that are indicated in the power spectra of the diffuse and specular components will be similar to each other. This suggests that specularities could provide reliable motion parallax information to a moving observer.

Keywords: motion, specularities, parallax, Fourier analysis

1. INTRODUCTION

In recent years, numerous computational methods have been developed for analyzing visual motion. Most of these methods are concerned with real-scene features, namely surface markings that remain fixed to the surface when the camera moves. The motion of virtual scene features, namely specular reflections, has received relatively less attention.* Nonetheless, specular reflections are highly prevalent in nature and, although they are not attached to a particular surface, their motion in many cases can give the scene its main visual character. It is difficult to imagine for instance a wavy water surface on a sunny day without considering the motion of the specular reflections. Indeed psychophysical experiments have shown that the presence of specularities helps the viewer to perceive surface shape.¹

Several computer vision methods have been developed to analyze the motion of specularities by tracking individual features over time.²⁻⁷ The goals there are typically to discriminate real features from virtual ones and to estimate properties of the underlying surface shape. These shape properties can be qualitative such as the sign of curvature in the case of an elliptical patch, or more quantitative such as a full 3D recovery of the surface profile. In this paper we consider the motion of a *field* of specularities over a random smooth surface, as seen by a laterally moving observer. For such a surface, the motion of specularities is complex and occurs in many different directions, not merely parallel to the observer motion *i.e.* in the epipolar direction.⁸ In particular, the caustic surfaces (see below) defined by the specularities are both in front and behind the reflecting surface. This suggests there is motion parallax between the specularities. We are interested in whether the motion parallax between the specularities resembles the motion parallax that occurs in any 3D scene where surfaces exist over a range of depths and the observer moves laterally. We address this question using the theory of motion parallax developed recently for the case of cluttered 3D scenes. This theory considers motion in the frequency domain.⁹

*Here we use the term specular reflection to refer to a mirror-like reflection off a shiny surface, such as water, ceramic, metal, glass, etc.

2. BACKGROUND

2.1. Motion of Specularities

Two general approaches have been taken to modeling the motion of specular reflections induced by a moving observer. Both approaches are based on the same underlying phenomenon. A 2D manifold of rays[†] arrives at a smooth surface and is reflected in a mirror-like fashion, i.e. each ray’s angle of incidence is equal to its angle of reflection. This results in a 2D manifold of reflected rays. An observer moves along a space curve, and traces out a 1D sub-manifold of these reflected rays. The direction of each of the rays that is observed defines the position of the specularity in the observer’s viewing sphere (image).

The first approach^{2,5,10–12} to modeling the motion of specularities is to express each specularity as a point on the mirror surface, namely the point of reflection. As mentioned above, the path of the moving observer intercepts a 1D sub-manifold of reflected rays. The motion of the specularity can be described as the path *on the mirror surface* where these rays begin, i.e. the points of reflection. The second approach^{6,7} considers the same 1D submanifold of rays that are observed, but now considers the envelope of this set of rays, namely the curve that is tangent to this 1D sub-manifold. This envelope is known as a *caustic*. As the observer moves, he intercepts different reflected rays and traces out different tangent points on the caustic. One can then speak of the specularity as moving along the caustic. The two approaches both describe the same set of rays, namely those that are mirror-reflected from the surface and then are observed by the moving observer. The difference between the two approaches is how the rays are parametrized. In the first approach, the rays are parametrized by where they intersect the mirror surface. In the second approach, they are parametrized by the tangent points on the caustic.




Figure 1. The caustic surface for a curved concave reflector. The path of the specularity can be considered either along the mirror surface or along the caustic.

Here we briefly review the main equation obtained using the first of these two approaches – namely, Blake’s specularity equation² – and we summarize a few observations that have been made. Similar observations can be made using the second approach as well. For simplicity, we consider only the case that the point light source is infinitely far away, so that the 2D manifold of incident rays is collimated. Let (u, v, n) be a coordinate system whose origin is some point p_0 on the surface where a specular reflection exists. The coordinates u and v define perpendicular directions within the tangent plane at p_0 , and the n coordinate defines the perpendicular distance from this tangent plane.

To say that p_0 is a specular point means that the surface normal bisects the angle between the light source and viewing directions. Let θ be the angle of incidence which is equal to the angle of reflection. The u and v

[†]The incident 2D manifold of rays can be collimated (point source at infinity) or diverging (from a point source at a finite distance). We consider only the collimated case in this paper.

directions within the tangent plane are chosen such that the direction of the light source is $(\sin\theta, 0, \cos\theta)$. By mirror reflection symmetry, the direction to the camera is $(-\sin\theta, 0, \cos\theta)$. Letting the distance to the camera be C , the position of the camera is $\mathbf{C} = C(-\sin\theta, 0, \cos\theta)$.

With these coordinates established, Blake showed that the instantaneous velocity (V_u, V_v) of the specularity, which is in the surface tangent plane, is related to the instantaneous 3D translation velocity (T_u, T_v, T_n) of the camera as follows:

$$\left(2C \begin{bmatrix} \frac{1}{\cos\theta} & 0 \\ 0 & \cos\theta \end{bmatrix} \begin{bmatrix} \frac{\partial^2 z}{\partial u^2} & \frac{\partial^2 z}{\partial u \partial v} \\ \frac{\partial^2 z}{\partial u \partial v} & \frac{\partial^2 z}{\partial v^2} \end{bmatrix} + \mathbf{I} \right) \begin{bmatrix} V_u \\ V_v \end{bmatrix} = \begin{bmatrix} T_u + T_n \tan\theta \\ T_v \end{bmatrix} \quad (1)$$

where \mathbf{I} is a 2×2 identity matrix.

Several interesting cases follow. First, if the surface curvatures are both high and the camera distance C is large, then (V_u, V_v) will be near zero, *i.e.* the specularities will move very little on the surface. A similar observation had been made earlier by Koenderink and van Doorn,¹¹ namely that specularities cling to regions of high curvature.

A second interesting case occurs if the velocity of the specularity (V_u, V_v) is very large. This can occur when the 2×2 matrix on the left side of the above equation is nearly singular. An example is that θ is 0 and the curvature[‡] of the surface in the direction (V_u, V_v) is roughly $-C$. This occurs when the surface is concave in this direction and the camera lies near the focal distance defined by this concavity.

A third interesting case occurs if the surface curvature is zero (*i.e.* all second derivatives of $z(u, v)$ vanish) and $\theta \approx 0$. In this case,

$$(V_u, V_v) \approx (T_u, T_v)$$

and so the specularity moves along the surface with roughly the same velocity as the observer. An example of this case is the reflection of the moon off still water.¹³

Later in this paper, we will see examples of these cases at work. For now, we turn to the method that we will use to analyze these motions.

2.2. Frequency analysis of motion

Blake's specularity equation is useful for understanding the *local* properties of specular motion. In this paper, we focus on the *global* properties of the motion, namely second order statistics of the time-varying image intensities. This global analysis is carried out in the frequency domain. Here we briefly review the theory on which our analysis is based.

Many studies have been conducted on motion in the spatiotemporal frequency domain. Watson and Ahumada¹⁴ introduced the notion of the motion plane property. They observed that if an image frame translates over time with a given velocity (v_x, v_y) , that is,

$$I(x, y, t) = I(x - v_x t, y - v_y t, 0)$$

then every component of the 2-D Fourier transform of the frame travels with the same velocity, and so

$$v_x f_x + v_y f_y + f_t = 0. \quad (2)$$

Here f_t and f_x, f_y are respectively the temporal and spatial Fourier coefficients. We refer to Eq. (2) as the *motion plane* model.

Langer and Mann⁹ extended this model to motion parallax, by introducing a model called *optical snow* in which there exists a set of parallel velocities $(\alpha v_x, \alpha v_y)$ within an image region, where α can take multiple values. An example is what a camera sees when moving laterally relative to a cluttered 3D scene.[§] (Also see the paper

[‡]Normal surface curvature in any direction is inversely related to the second derivative of z in that direction.

[§]This could be a moving camera in a cluttered scene such as the forest, or a fixed camera looking at falling snow.

by Rivait and Langer in these Proceedings.) The optical snow model implies that a family of motion planes exists in the 3D frequency domain, i.e.

$$\alpha v_x f_x + \alpha v_y f_y + f_t = 0 .$$

Since all of these planes pass through a common axis, the family of planes has the appearance of a *bowtie*. In cluttered scenes, there are many distinct surfaces present, and these give rise to a discrete set of motion planes. The optical snow model does not require, however, that many distinct surfaces are present. It merely requires that surfaces are present at many different scene depths. The model would also apply to a continuous “ground plane” surface that is slanted in depth. In this case, a continuum of motion planes would result. See Fig. 2.

An central assumption of the optical snow model is that the surfaces are Lambertian, and so the image features correspond to the projection of scene features that are *fixed to the surface*. This is quite different from the case of specularities which move along the surface, as discussed in Sec. 2.1.

For Lambertian surfaces, the image velocities are as follows. Let $z(x, y)$ represent the depth of a surface point that is visible at position (x, y) in the image plane. Assume the observer moves laterally in the x -direction with speed T_x , and the image plane is at $z = 1$. Then the image speed is inversely proportional to depth,

$$(v_x, v_y) = \left(\frac{-T_x}{z(x, y)}, 0 \right).$$

Note that the α values for the optical snow model are essentially determined by the inverse depths $1/z(x, y)$ since T_x is a constant. For a continuous surface with a range of depths, a continuum of motion planes is obtained. For the cases that we are assuming, each of these motion planes intersects the f_y axis. Thus, if we project the 3D power spectrum by summing the power in the f_y direction, then we obtain a 2D bowtie pattern in (f_x, f_t) space. See Fig. 2. Examples for rendered sequences will be given later in this paper.

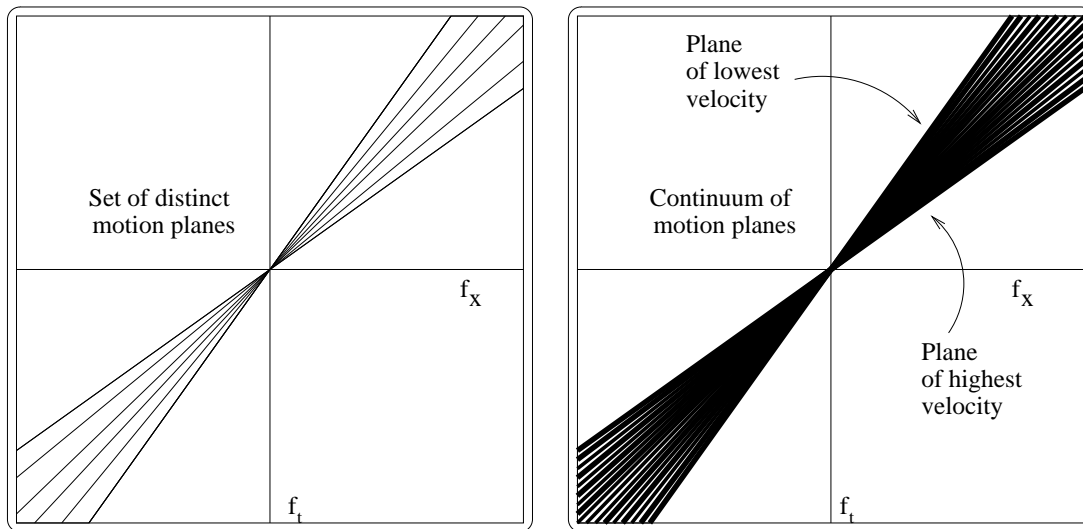


Figure 2. (a) The projected 2D bowtie pattern for a scene with many depth discontinuities. (b) The projected 2D bowtie pattern of a scene with a continuous depth map (smooth undulating surface).

2.3. Specularities and motion parallax

The main question we address in this paper is whether the optical snow and bowtie models, which apply for Lambertian (diffuse) scenes, also apply to scenes in which the surface reflectance has a specular component. One might expect that the motion parallax from specularities behaves very similarly to optical snow. A mirror surface focuses the incident light, defining caustic surfaces both in front of and behind the mirror. If the surface

has many local maxima and minima, then many distinct caustic surfaces should result. Because these caustic surfaces have points that lie at a range of depths, one might expect such points to give rise to motion parallax.

However, strictly speaking, this hypothesis is not quite correct. First, specularities can appear and disappear[¶] as the observer moves and this is not addressed by the hypothesis. Second, specularities move along the caustics as the observer moves. This is quite different from the Lambertian case, where features are fixed to the surface. Given these two issues, we were curious to know whether a weaker form of the hypothesis holds nonetheless.

Our experiments consider the case of low-relief height field (terrain) surfaces.^{||} We study how the motion of specularities varies with the parameters of the height field, and how this motion compares to that of the Lambertian component. In particular, the observer moves in the x direction relative to a surface $z(x, y)$ and we examine the projected bowties in the power spectra of the specular and diffuse (Lambertian) components of the image sequences.

3. EXPERIMENTAL RESULTS

3.1. Surfaces and image sequences

Our experiments used surfaces that were rendered using OpenGL. Each height field surface is defined by a sum of 2D sinusoids. Each surface is a triangular mesh defined on a square grid of size 1028×1028 . Each surface is generated using the spectral synthesis method. A white noise height field is lowpass filtered with a sharp cutoff at $k < k_{max}$ where $k = \sqrt{k_x^2 + k_y^2}$ and $k_{max} = 42$ cycles. The surface height is scaled so that its standard deviation is one of three values: 0.1, 0.5, or 2.0. We will refer to these as the *relief values* of the surfaces. Note that in each case the surfaces are low relief, since each standard deviation is very small (order 1) in comparison to the width of the surface (order 200, see below).

Two scenarios were considered – see Fig. 3. The first is highly prevalent in nature, namely a low relief shiny ground plane surface that is viewed from an oblique angle. In this scenario, there is a large range of depths in both the real and virtual (caustic) surfaces. In particular, the lower part of the image is much closer to the camera than the upper part. The second scenario takes the same surface but views it from a frontoparallel direction. This condition serves as a control, where the range of real surface depths z relative to the camera is very small across the image.




Figure 3. (a) Scenario 1: Slanted ground plane. The scene exhibits a wide range of depths. Points in the lower of the visual field are much closer to the camera than points in the upper visual field. (b) Scenario 2: Frontoparallel surface. The scene exhibits a narrow range of depths.

[¶]These appearance/disappearance events happen in pairs when the observer crosses a caustic surface.^{5,10,11}

^{||}Low relief means that the height variations are small in comparison to the surface width.

In both scenarios, the surface has depth centered at $z_0 = 200$ at the middle row of the image. In the first scenario, the surface is slanted by 45 degrees with respect to the x -axis so that points in the bottom row of the image have depth z much less than 200 and points at the top row of the image have depth z much greater than 200. In the second scenario, all points have roughly the same depth (200).

The field of view was 53 degrees. The observer moves in the x direction, with speed T_x that is conveniently chosen such that a point at depth $z_0 = 200$ travels exactly one image width in the $T = 128$ frame sequence. Such a point has a speed $v_x = 1$ pixels per frame, and this corresponds to a motion plane of slope 1.

A different collimated light source was chosen for the two scenarios. (See Fig. 3.) In the slanted surface scenario, the incident light is from above, that is, in the camera's y -axis direction. In the frontoparallel scenario, the incident light is parallel to the line of sight. These lighting conditions were chosen so that the specularities were distributed roughly symmetrically about the image center in both scenarios.

Separate color channels (red and blue) were attributed to the Lambertian and specular components, in order to separate them for the purpose of comparison. For the specular component, the OpenGL *shininess* parameter was set to 128. Notice that this does not define a perfect mirror, unlike in Eq. (1).

Each sequence then consisted of 128 image frames each of size 128×128 pixels. Prior to computing the 3D Fourier transform $\hat{I}(f_x, f_y, f_t)$ of each sequence, the rendered image sequence was multiplied by a Hanning window in each of the x, y, t dimensions. The power spectrum $|\hat{I}(f_x, f_y, f_t)|^2$ of the (windowed) image sequence was then summed in the f_y direction, to yield a function of variables f_x and f_t . Power spectrum plots below show the average over 20 sequences.




Figure 4. One frame of the diffuse component for each of three surface reliefs. (See text for discussion.)

3.2. Scenario 1: Slanted ground plane surface

One frame of each of the three types of sequences is shown in Fig. 4 (diffuse component) and Fig. 5 (specular component). The three columns in each figure show the results for the three surface reliefs: 2, 0.5, 0.1.

The wide range of image speeds in the diffuse components (Fig. 4) are due to the global surface slant, namely points at the top of the image are further away and thus moving more slowly in the image. The local height variations contributed only little to the range of image speeds since the surfaces are of such low relief. Although the diffuse component (Fig. 4) have roughly the same range of image speeds for each of three surface reliefs, the distribution of power varies for the three relief values, namely more power is present at high frequencies when the relief value is greater. These power spectrum differences are due to differences in surface shading.

Figure 5 shows corresponding examples for the specular components. In both (a) and (b), a well-defined bowtie pattern exists as in the diffuse case, though there clearly is more power present at higher frequencies in the specular case. This increased high frequency power is not surprising, since the intensity varies more quickly with the surface normal in the specular case than in the diffuse case. However, apart from this increased power at high frequencies, the bowties seem quite similar at first glance in the specular vs. diffuse cases.




Figure 5. The same frames are shown as in Fig. 4, but now for the specular component.

A closer examination of the bowties indicates other interesting differences between the diffuse and specular cases. These differences concern the range of speeds present. At high speeds (high motion plane slopes) and high spatial frequencies, there is relatively more power in the specular component than in the diffuse component (see Fig. 6). We believe there are two reasons for it.

First, note the high speed motion planes correspond to nearby parts of the surfaces and these parts tend to contribute relatively more power at lower spatial frequencies. This is visually evident in the single image frames:

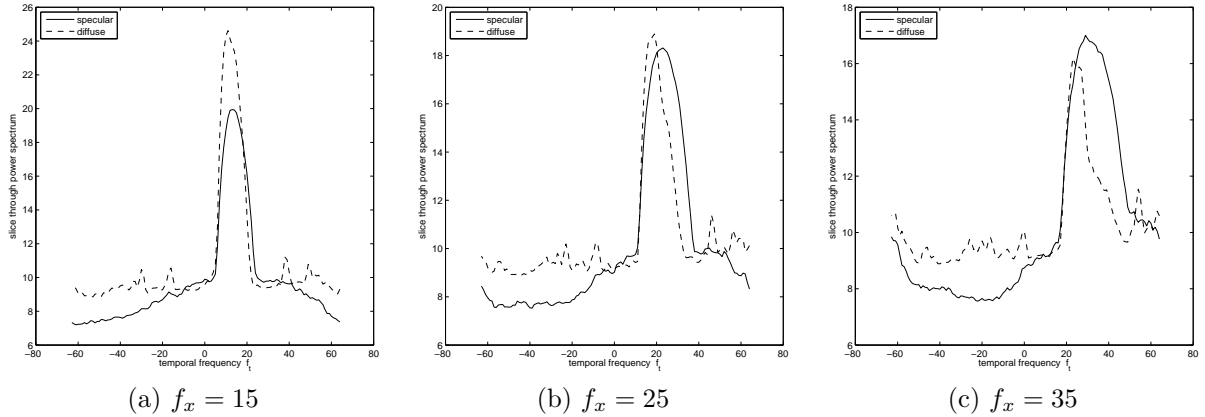


Figure 6. A plot of power vs. temporal frequency f_t for spatial frequency slices $f_x = 10, 25, 35$ from Figs. 5 and 4. Whereas in (a), the two curves are aligned, in (b,c) there is relatively more power in the specular component (see text).

the lower parts of the specular and diffuse images of Figs. 4 and 5 (a,b) are more dominated by lower spatial frequency components. This dominance is less severe for the specular case though, since the specularities have relatively sharp intensity transitions which do contribute some high spatial frequency power.

A second reason for the relative lack of power in the diffuse power spectrum at high spatial frequencies and high speeds can be understood by recalling Eq. (1). All other variables being equal, the speed of specularity along the surface is inversely related to the distance C to the surface. In particular, since the statistics of the surface curvature are indeed the same over the entire surface, the speeds of the specularities on the surface should be smaller for surface points that are at greater depth. (We emphasize that this is a statement about the speeds *on the surface*.) As $z \rightarrow \infty$, the speeds of the specularities on the surface goes to zero. It follows that the image velocities of the diffuse and specular components should be very similar for points that are far away. This constraint does not hold, however, for points that are near to the camera. For these points, the motion of the specularities on the surface can be quite large, and hence there can be a greater range of motions in the specular component than in the diffuse component. It follows that, at high speeds, the specular components should have a greater range of velocities. We would expect this effect to be evident in all three plots of Fig. 6. The relatively limited effect observed in Fig. 6(a) could be due to the inherent limitation of speed resolution in the power spectrum at lower frequencies.

Finally, we briefly mention the lowest relief case namely Figs. 4,5(c). In this case, the curvature of the surface is smallest and so the motion of the specularities on the surface can be relatively large. This motion of the specularity can be in any direction and can yield a wide range of image velocity directions. Even though there is more spreading of power in (c) than in (a,b), however, the motion of the specular component is still predominantly in the direction of the diffuse component. This is because the specular component is the sum of the diffuse component and the component due to motion of the specularity on the surface.

3.3. Scenario 2: Frontoparallel surface

Finally, we address the second scenario (recall Fig. 3(b)), in which there is only a small range of depths z in the scene. The diffuse component induces almost no motion parallax and the bowtie collapses to a single motion plane (of slope 1). See Fig. 7. There is a slight spreading in the power spectrum about the line of slope 1, but this is mainly due to the Hanning window.

The data for the specular component is shown in Fig. 8. Once again, in (a,b) the specular components display a similar signature as the Lambertian components, namely most of the power is concentrated along the line of slope 1. In (c), the motion is predominantly near the plane of slope 1 but as in the slanted plane scenario there is significant deviation from this average motion. This deviation is due to the low surface curvature, which permits the specularities to move more quickly on the surface. Since the direction of the specular motion need not be parallel to the motion of the diffuse component, a non-bowtie spreading of power can occur.




Figure 7. One frame is shown for the frontoparallel Lambertian surface, and for the case of relief value 2. No motion parallax occurs.




Figure 8. Single frames for the specular component in the frontoparallel case. Projected power spectra are concentrated along the line of slope 1.

4. CONCLUSION

Previous work on the motion of specularities has concentrated on a local analysis of the motion. There, the goals were typically to try to distinguish the diffuse and specular components of an image sequence or to reconstruct the local surface shape. Our goals and approach are different. We have investigated the global field of specularities, and have asked to what extent the range of image velocities present in the specular component is similar to the range present in the diffuse component. We found that the range of image velocities is indeed similar, if the surface relief is low but not too low. In particular, we showed that in the case of the lateral observer motion and low relief surfaces, the specular component exhibits broadly the same spatio-temporal power spectrum (same direction and velocities) as the diffuse component. Some discrepancy was found for nearby parts of the surface, which produce both high speeds in the diffuse component and high speeds in the specular motion along the surface. One interesting question for future research is whether this discrepancy between the motions of the diffuse and specular components is large enough to be detected by the human visual system, and if so, whether it is useful or harmful for perceptually discriminating and decomposing surfaces into their diffuse and specular components.

Acknowledgements

This research was supported by an NSERC Discovery Grant to M. Langer.

REFERENCES

1. J. F. Norman, J. T. Todd, and G. A. Orban, "Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information," *Psychological Science* **15**, pp. 565–570, 2004.
2. A. Blake and G. Brelstaff, "Specular stereo," in *International Joint Conference on Artificial Intelligence*, pp. 973–976, 1985.
3. H. Schultz, "Retrieving shape information from multiple images of a specular surface," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**(2), pp. 195–201, 1994.
4. A. C. Sanderson, L. E. Weiss, and S. K. Nayar, "Structured highlight inspection of specular surfaces," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(1), pp. 44–55, 1988.
5. A. Zisserman, P. GIBLIN, and A. Blake, "The information available to a moving observer from specularities," *Image and Vision Computing* **7**, pp. 287 – 291, 1989.
6. M. Oren and S. K. Nayar, "A theory of specular surface geometry," *International Journal of Computer Vision* **24**(2), pp. 105–124, 1996.
7. R. Swaminathan, S. Kang, R. Szeliski, A. Criminisi, and S. Nayar, "On the Motion and Appearance of Specularities in Image Sequences," in *European Conference on Computer Vision (ECCV)*, **I**, pp. 508–523, May 2002.
8. R. C. Bolles, H. H. Baker, and D. H. Marimont, "Epipolar-plane image analysis: An approach to determining structure from motion," *International Journal of Computer Vision* **1**, pp. 7–55, 1987.
9. M. S. Langer and R. Mann, "Optical snow," *International Journal of Computer Vision* **55**(1), pp. 55–71, 2003.
10. M. S. Longuet-Higgins, "Reflection and refraction at a random moving surface. i. pattern and paths of specular points," *Journal of the Optical Society of America* **50**(9), pp. 838–844, 1960.
11. J. J. Koenderink and A. J. van Doorn, "Photometric invariants related to solid shapes," *Optica Acta* **27**(7), pp. 981–996, 1980.
12. S. Waldon and C. R. Dyer, "Dynamic shading, motion parallax and qualitative shape," in *Proceedings of the IEEE Workshop on Qualitative Vision*, pp. 61–71, 1993.
13. A. Blake and H. H. Bulthoff, "Shape from specularities: computation and psychophysics," *Phil. Trans. Roy. Soc. B* **331**, pp. 237–252, 1991.
14. A. Watson and A. Ahumada, "Model of human visual-motion sensing," *Journal of the Optical Society of America* **2**(2), pp. 322–342, 1985.