Shape from Specularities

Andrew R. Banks

Introduction

This page commences by giving an overview of what specular reflections are, followed by an account on their use to aid three-dimensional shape estimation.

Specular Reflections

For a specular reflection, the reflection of a surface is stronger in one direction than another, resulting in a *specular highlight*. Specular highlights are visible on smooth, shiny surfaces, for example polished metal. This type of reflection contrasts ambient and diffuse reflections since they are not dependent upon the viewpoint. Figure 1 shows a photograph of a car window with a clearly visible specular highlight.

Figure 1: Specular highlight upon a car window [5]

Figure 2 shows a cross-section of a curved surface, where **N** is the surface normal, **L** the light source direction and **R** the specular reflection direction at its greatest intensity.

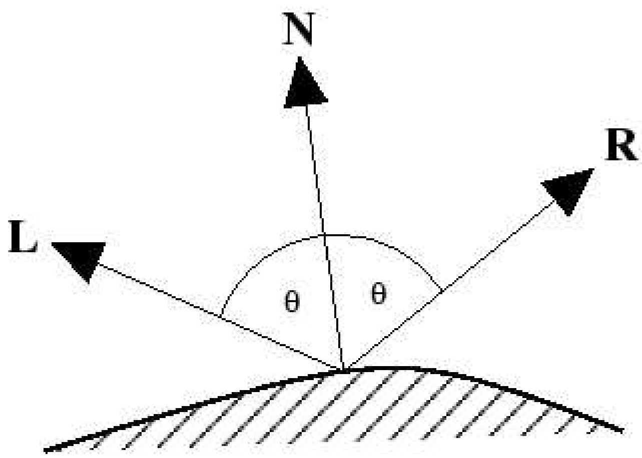

Figure 2: The condition for specular reflection

For a perfect mirror, the angle of incidence θ equals the angle of reflection. Now, if **R** and the viewpoint direction, **V**, coincide, then the specular reflectance is visible. It is worth noting that in reality, these two vectors do not have to coincide exactly: **V** is permitted to reside within a range of values constrained by the shininess of the surface. The shinier the surface, the more closely **V** must approximate **R**.

Specular Reflections and Surface Shape

There are many methods of reconstructing a three-dimensional shape from an image, including the use of textures, shadows and shading. Indeed the human visual system uses a combination of these methods in shape estimation. However, most research in shape estimation from specular reflections is undertaken using perfect mirrors, isolating specular reflections from these other visual cues [1].

Figure 3 shows three images of a computer generated three-dimensional shape that is a perfect mirror. Each image shown has been captured when the object has been placed in a different scene.

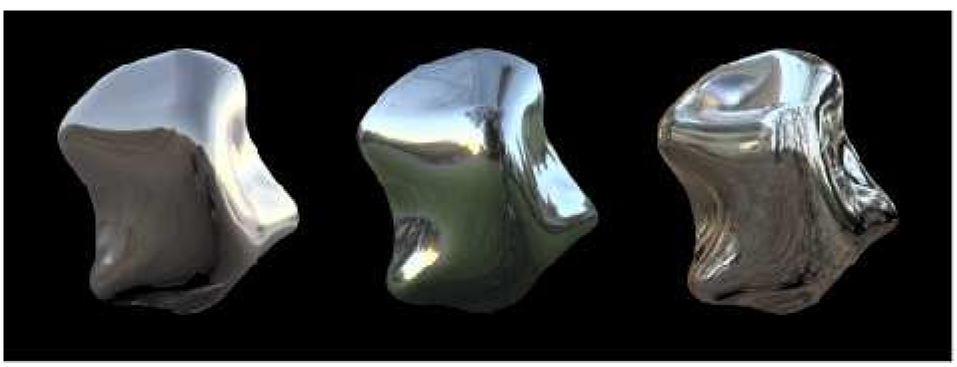

Figure 3: Snapshots of a specular object placed in distinct environments [2]

Evidently the appearance of the object is greatly dependent upon the environment in which it resides. It is clearly seen that the reflections are highly unstable; it is a non-trivial problem to perform recognition on environmental objects through these reflections. On the other hand, the appearance of the shape itself is relatively stable across the three images.

Main Approaches

There are two main approaches taken to solving this problem, as identified by Fleming et al. [2]. The first approach requires that the system has knowledge of the scene geometry. It performs by recognising a scene object from its distorted reflection upon the specular surface, and then estimating the transformation in which the surface has applied to it. This approach has two disadvantages; it is not only computationally expensive, but also requires detailed knowledge about the environment.

The second approach is to assume the world is a texture, one in which remains stable across difference scenes, including the distribution of orientations is maintained over the scenes. The surface distorts this *texture* according to its shape, producing patterns similar to textures across its surface; these are termed *orientation fields*.

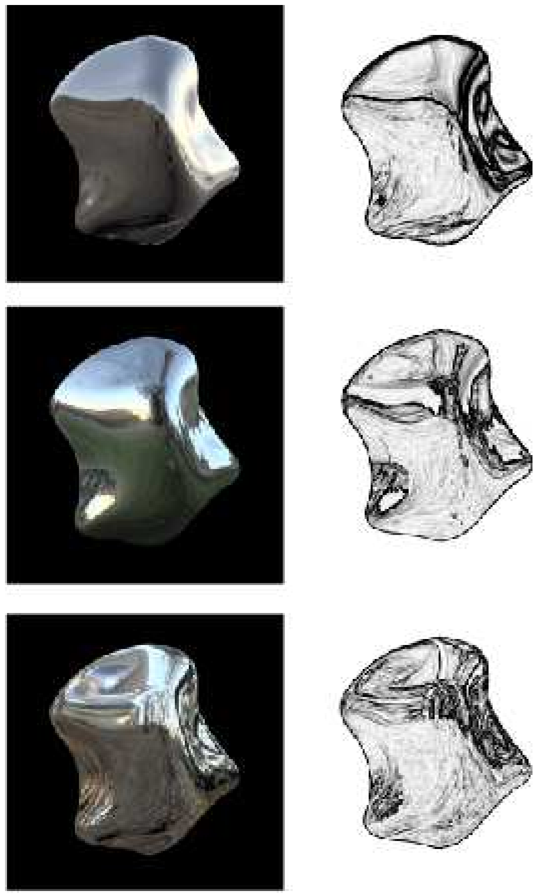

Figure 4: Snapshots of a specular object placed in distinct environments: Left-hand colum shows the object. Right-hand column shows the object after simple edge-detection [2]

The left-hand column of Figure 4 shows the appearance of the specular object taken earlier placed in three different scenes (the context has been cropped out for clarity). The right-hand column shows the output of each image after some simple edge-detection was performed upon it. It can be seen quite clearly that the distortion fields remain relatively consistant over the three scenes.

A system which uses distortion fields in the reconstruction of a three-dimensional shape clearly has an advantage over the first method identified, such that no knowledge of the scene gemoetry is required.

The authors show that by taking the world reflection on the object as a texture, it is not correct to use the same method of shape estimation as *real* textures with specular objects (at least not on perfect mirrors). This is so, because with textures, the compression of the image is dependent upon the slope of the surface, that is the first surface derivative. However, with specular reflections, the compression is dependent upon the rate at which the surface normals change, in other words, the second surface derivative.

Related Reading

In addition to works already discussed, the reader may also find the following resources useful.

- Solem et al. [5]. The authors exploit the condition of specular reflections, namely the constraining of surface normals, along with a smoothness condition in the reconstruction of a specular surface from an image sequence. The images are interestingly captured from a standard hand-held video camera
- Halstead et al. [3]. This paper applies the use of shape from specularities to the measurement of the human cornea, in which an algorithm is presented
- Savarese et al. [4]. The authors focus upon the first approach described earlier in Section 3.1, in which knowledge of the scene is exploited for shape reconstruction
- Zheng et al. [6]. In this paper, the authors use the motion of specular highlights during the rotation of specular objects in order to reconstruct the object

Bibliography

1
: FLEMING, R. W., AND BÜLTHOFF, H. H.
Orientation fields in the perception of 3D shape.
Journal of Vision 5, 8 (September 2005), 525-525.

2
: FLEMING, R. W., TORRALBA, A., AND ADELSON, E. H.
Specular reflections and the perception of shape.
Journal of Vision 4, 9 (September 2004), 798-820.

3
: HALSTEAD, M. A., BARSKY, B. A., KLEIN, S. A., AND MANDELL, R. B.
Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals.
Proceedings of the ACM SIGGRAPH Conference on Computer Graphics (1996), 335 - 342.

4
: SAVARESE, S., AND PERONA, P.
Local analysis for 3D reconstruction of specular surfaces.
Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001 vol.2 (2001), 738 - 45.

5
: SOLEM, J. E., AND HEYDEN, A.
Surface shape from specularities.
Lecture Notes in Computer Science (2003), 828-835.

6
: ZHENG, J. Y., AND MURATA, A.
Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources.
IEEE Transactions on Pattern Analysis and Machine Intelligence 22, 8 (2000), 913 - 920.

Footnotes

... cues [1]
: In contrast to this, Fleming et al. [1] explore commonalities between such visual cues

---

[Andrew R. Banks](#) (s0568337)
Fri Feb 3 2006