



Specular Surface Recovery via Coded Target Stereopsis

Arlene Ripsman

Michael R. M. Jenkin

Piotr Jasiobedzki

Technical Report CS-2003-03

March 26, 2003

Department of Computer Science

4700 Keele Street North York, Ontario M3J 1P3 Canada

Specular Surface Recovery via Coded Target Stereopsis

*A. Ripsman*¹, *M. Jenkin*¹, *P. Jasiobedzki*²

¹Department of Computer Science and Centre for Vision Research
York University, Toronto, Ontario, Canada M3J 1P3

²MD Robotics

Brampton, Ontario, Canada L6S 4J3

{arlene,jenkin}@cs.yorku.ca, PJASIOBE@mdrobotics.ca

Abstract

A key challenge facing space-based computer vision systems is the presence of highly specular surfaces. Such surfaces present a challenge to conventional vision systems due to specular reflections, which may mask the true location of objects and lead to incorrect measurements. The problem is complicated by the fact that the use of high-powered illuminants, such as laser beams, can be problematic since the instruments inside space structures are often sensitive to various forms of radiation.

This article presents a new approach to reconstructing the local surface structure of highly specular objects. The system utilizes a commercial trinocular stereo vision system and a coded low-power two-dimensional illuminant to recover surface structure. The local surface structure of an object is obtained by projecting light patterns onto the object's surface and inferring surface structure by examining the interaction of the illuminant and the surface. The system recovers both specular and diffuse surface regions. Experimental results demonstrate the effectiveness of the approach.

1 Introduction

Objects are being sent into space with increasing frequency and it is important to find ways to extend their working lifetime. It is generally difficult, and in some

cases impossible, to retrieve such objects for regular inspection and repair due to their large size. As a result, the *in-situ* inspection and repair of these types of objects is required.



Figure 1: The Hubble Space Telescope, NASA photograph.

Man-made space objects are typically covered with a metallized film that renders the surface of such objects highly specular. This specular surface is used as a reflective insulator to protect both human occupants and sensitive on-board instruments from solar radiation. Figure 1 is a NASA photograph of the Hubble Space Telescope on orbit and illustrates the nature of this reflective insulator. The reflection of the astronaut is clearly visible reflected in the telescope's surface. Traditional vision-based sensing systems are designed with non-specular surfaces in mind. In order to inspect highly specular surfaces visually, new algorithms and sensing systems are required.

In addition to their specular surfaces, orbital objects present a number of unique problems for inspection systems. The instruments on board orbital objects

may be sensitive to active illuminants such as lasers, radar and similar detecting devices. As a result, inspection systems that utilize no extra illuminant or a low-powered illuminant are preferable to high-power systems. Finally, the surfaces of orbital objects are a combination of highly specular and highly diffuse surfaces, and thus, in order for the system to be of some practical use, it must be able to recover the structure from both types of surfaces.

This article describes a novel approach to the reconstruction of surfaces of highly specular objects. Rather than attempting to recover the surface structure of a specular object directly, we approach the problem by looking at the effects of specular surfaces on the appearance of other objects reflected within them. Specifically we utilize a commercial trinocular stereo vision system, coupled with an addressable video target to project two-dimensional light patterns, in order to recover three-dimensional points on the surface of highly specular objects.

2 Previous Work

Surface reconstruction is a huge and ongoing area of research in computer vision. Traditional approaches include binocular stereopsis (described in detail in [19]), multiple-baseline stereo [16], coded structured light (see [35] for a review), depth from focus/defocus [52] interferometry [3], shape from shading [10], photometric stereo [49, 50], shape from texture [6, 22, 5], shape from occluding contours [2], shape from shadows [39], shape from motion [11, 18] and shape from polarization

[20, 26, 48] to name but a few approaches. Research in surface reconstruction typically assumes that the surface of the objects are Lambertian. That being said, there has been some research in extending several of the methods to include specular surfaces. In particular, photometric stereo and shape from shading has been used by several researchers to recover specular surfaces [12, 1, 25, 27, 36, 41] and more recently the task has been formulated for artificial neural networks [4, 13, 32]. Other research methods for specular surfaces include shape from specular motion [21, 34, 53, 54, 55], polarization methods [20, 26, 30, 31, 46, 47, 48], structured light for specular surfaces [8, 9, 17, 38, 42] and shape from color [37, 43].

Each of the techniques used to recover specular surfaces have limitations. Many assume that the surfaces are glossy rather than highly specular, have uniform reflectance, or the techniques are meant to be applied to small objects at a known location. Most of the techniques produce only shape information rather than range data. Reconstructing the surface of orbital objects present unique challenges not typically encountered in other reconstruction applications. At present there does not exist an algorithm capable of recovering the three-dimensional surface structure of highly specular orbital objects.

3 Specular Surface Geometry

Specular surfaces are those surfaces that exhibit specular reflections. A specular reflection is defined as:

“Reflection of electromagnetic, acoustic or water waves in which the reflected waves make equal angles with a line perpendicular to the reflecting surface, and lie in the same plane with it.”¹

The definition of specular reflectance is commonly known as the law of reflection. For non-specular or diffuse surfaces, incident light is scattered by the surface so that the perceived brightness is independent of the direction of view. Unlike diffuse surfaces, perfectly specular surfaces reflect all the light in one direction. One important result of this is that an object reflected in a planar mirror appears behind the surface of the mirror by exactly the same distance that the actual object is in front of the mirror (see [7] and also the proof that follows). The properties of specular reflectance are used by the Passive Target Technique, which forms the basis for the algorithm described in this paper.

4 Passive Target Technique

The Passive Target Technique (see [14]) uses specular surface geometry to infer points on the surface of a perfectly specular planar object. If a known illuminated

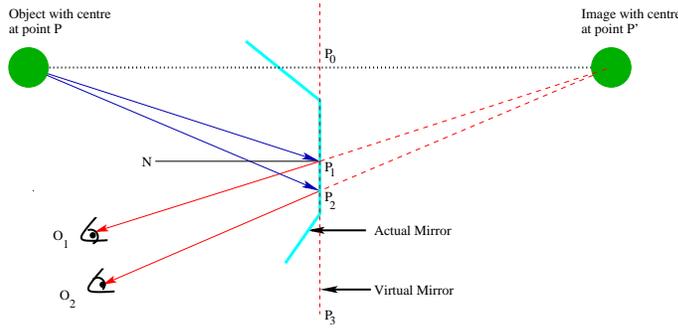
¹*From the McGraw-Hill Dictionary of Scientific Terms, Fifth Edition, page 1882.*

target point, P , is projected onto a specular surface, its corresponding virtual point P' in the image is the reflection of P . Knowledge of P and P' can be used to infer points on the surface of the mirror. Suppose that P' is viewed and recovered by a system with multiple viewers. For a perfectly specular planar surface (a mirror), the real distance from the object to the mirror and the virtual distance from the mirror to the image are always equal (see Figure 2). Using this principle, together with the locations of P and P' , we calculate a point P_0 on the virtual planar mirror in which P' is reflected. P_0 is calculated as

$$P_0 = (P + P')/2 \tag{1}$$

The virtual planar mirror in which P is reflected is then calculated as the plane passing through the point P_0 that is perpendicular to the line segment connecting points P and P' . The intersection of the virtual mirror with a ray joining the observer to the virtual point P' will result in a point on the surface of the mirror. If P' is viewed by multiple cameras then each camera obtains a point on the mirror.

Finally, it should be noted that, as shown in Figure 2, the virtual mirror and the point P_0 on the virtual mirror do not need to exist. The mirror only needs to exist and to be planar in the region of the reflection of the rays joining the observers to P' . The passive target technique provides the basis for a method to recover points on the surface of highly specular objects. In order to use this



Let P_1 and P_2 be two points on a locally planar specular surface which reflect light from a source point P into two camera lenses O_1 and O_2 . If we draw a line segment from the source point P to its reflection P' , then we can define a virtual mirror on which P_1 and P_2 will lie such that $PP_0 = P_0P'$ and $PP_0P_1 = 90^\circ$.

Proof:

$$\begin{aligned}
& \angle PP_1N = \angle NP_1O_1 \\
& \therefore \angle PP_1P_0 = \angle O_1P_1P_2, \\
& \angle O_1P_1P_2 = \angle P'P_1P_0 \\
& \therefore \angle PP_1P_0 = \angle P'P_1P_0 \\
& \sim \angle PP_2P_1 = \angle P'P_2P_1 \\
& \therefore \triangle PP_2P_1 = \triangle P'P_2P_1 \\
& \quad PP_1 = P'P_1 \\
& \therefore \triangle PP_1P_0 = \triangle P'P_1P_0 \\
& \quad PP_0 = P'P_0 \\
& \therefore \angle PP_0P_1 = \angle P'P_0P_1 \\
& \quad \angle PP_0P_1 = 90^\circ
\end{aligned}$$

Therefore, $PP_0 = P'P_0$ and the plane defined by the mirror P_0P_3 is perpendicular to PP' .

Figure 2: Recovery of points on a specular surface using the passive target technique.

method certain issues must be addressed,

- **Determining P .** What sort of illuminant should be used and how should its position be determined?
- **Determining P' .** How can the virtual reflection of P be recovered?
- **Determining the correspondence between P and P' .** The key to the

passive target approach is to establish the correspondence between the real point P and its reflection P' . How can this correspondence be constructed?

- **Separating specular and diffuse points.** The passive target approach assumes perfectly specular surfaces. In practice, surfaces contain both specular and diffuse components. How to distinguish between specular and diffuse results?
- **Dealing with non-planar surfaces.** The passive target approach considers only perfect plane mirrors. How to deal with locally planar mirrors that are not globally planar?
- **Calibration.**

5 Coded Target Specular Stereo (CTSS)

Here we describe an algorithm that recovers specular surfaces using a trinocular stereo vision sensor coupled with a two-dimensional addressable illuminant. The underlying concept is to use the trinocular camera system to record three-dimensional locations (P') and to use the addressable two-dimensional illuminant to generate illuminants at known locations (P).

The CTSS algorithm is essentially a six step algorithm with which virtual points P and the corresponding illuminant points P' are generated and identified and combined to recover the surface structure. In brief, the algorithm can be

described in terms of the following steps:

1. The portion of the object that is to be reconstructed is illuminated. A purely specular surface will have no inherent texture of its own, and thus will be invisible unless some textured pattern is reflected within it. In order to break up this camouflage, an “unstructured” illuminant is used. The unstructured illuminant process is described in Section 5.2.
2. Three-dimensional points are obtained from the illumination of the surface. The Triclops system [29] is used to recover three-dimensional points (both specular and diffuse) within its field of view.
3. The recovered three-dimensional points are divided into specular (virtual) and diffuse (real) points. This process is described in Section 5.3.
4. For each specular point recovered, the true position of the illuminant that gave rise to the point is established. This is described in Section 5.4.
5. The Passive Target Technique is used to transform the virtual specular points into real points on the object’s surface.
6. The diffuse and specular points are merged into a single representation of the object.

The following subsections describe the system and provide details of the algorithm.

5.1 Experimental Setup

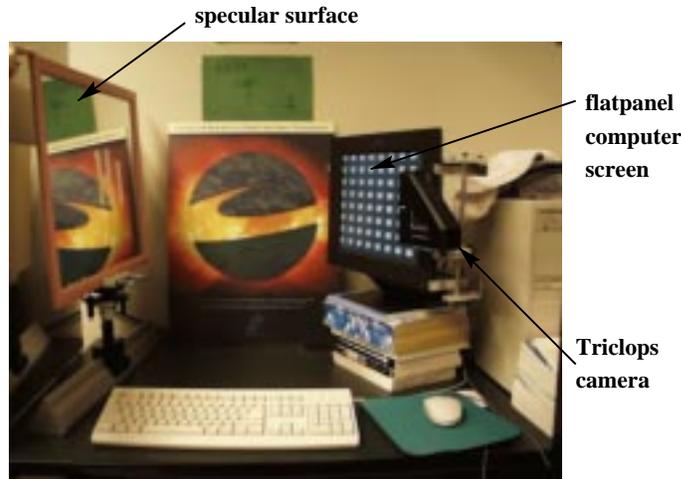


Figure 3: Experimental Setup. The sensor consists of a flatpanel computer screen to generate the illuminant and a Triclops trinocular camera module.

The laboratory setup, shown in Figure 3, is very simple. It consists of a Triclops camera module (see below) rigidly affixed to a flatpanel LCD computer monitor. The computer monitor provides the necessary controlled two-dimensional illumination.

5.1.1 Triclops Stereo Vision System

The Triclops Stereo Vision system [29] is a commercial vision system developed at the University of British Columbia and now marketed by Point Grey Research Inc. The system consists of a camera module and an accompanying software system. The camera module contains three $\frac{1}{3}$ " CCD cameras and simultaneously captures three gray-scale images of the scene. The Triclops software system uses a proprietary multi-baseline approach to stereo processing based on work of Kanade

et al. [15] and computes a disparity image from which the three-dimensional location of points in the scene can be determined.

5.1.2 Screen Calibration

In order to use the passive target technique it is necessary to calibrate the illuminant: For each point on the illuminant, it is necessary to determine the three-dimensional location of the point in the coordinate system defined by Triclops. This is accomplished by establishing an affine mapping between pixel locations and three-dimensional points on the illuminant. The illuminant calibration needs to be done only once and must be repeated only if the camera-illuminant configuration is changed. Details of the illuminant calibration process can be found in [33].

5.2 Unstructured Light and the Correspondence Problem

In order for Triclops to solve the correspondence problem it must have visual features, or texture, that can be correlated between the three images. Since most of the applications we are looking at tend to involve objects that are large, relatively smooth and featureless, we cannot guarantee that the surface will contain sufficient texture or that it will reflect sufficient naturally occurring texture. In order to avoid this problem we provide the required texture using light patterns that are projected onto the object. We project n (where $n > 2$) different patterns onto the object. From Triclops a disparity map is obtained for each pattern. For

each image location (pixel), if half or more of the disparities are within some threshold λ of one another, then the average of these disparities is used to arrive at a three-dimensional position. If the disparities differ by more than λ then results for this image location are not computed. We have achieved good results with $n = 5$ and $\lambda = 30$ pixels. The patterns are used to breakup camouflage and almost any sufficiently textured pattern would be suitable. The patterns used

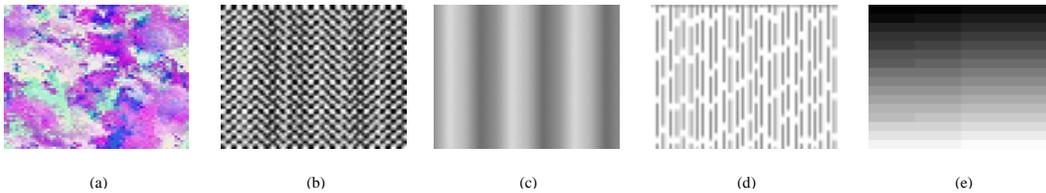


Figure 4: Unstructured light patterns used to produce the disparity maps.

currently are shown in Figure 4. The patterns include simple textures (patterns (a)[44] and (e)) as well as patterns proposed by other researchers. Pattern (b) was proposed by Vuylsteke et al. [45], pattern (c) was proposed by Wust and Capson [51], and pattern (d) was proposed by Maruyama et al. [23].

5.3 Specular and Diffuse Surfaces

In order to reconstruct highly specular planar surfaces with diffuse components, it is necessary to determine which image points correspond to specular reflections and which correspond to diffuse illumination. Once this has been accomplished, the three-dimensional location of diffuse points can be obtained directly from the disparity image produced by the Triclops system. The three-dimensional location

of specular points can be determined by the passive target method.

In order to divide the recovered points into specular and diffuse for a non-planar surface, an image is taken after each of the patterns described in Section 5.2 are projected onto the test object. For each pixel, i , the gray-scale intensity in each of Triclops' three images are recorded. We let $I_{i,max}$ and $I_{i,min}$ be the maximum and minimum intensities respectively for pixel i over all of the images. The system determines that a pixel, i , is diffuse iff

$$I_{i,max} - I_{i,min} < \sigma \quad (2)$$

Where σ depends on the allowable range of gray-scale values in the images. We use a range of 256 gray-scale values and σ is set to 60. The system assumes a point is diffuse when its intensity is relatively invariant to the light emitted by the projector, i.e. the difference between the minimum and maximum intensities is very small.

5.4 Structured light to identify and locate the illuminant

Given a disparity image, it is necessary to identify which point on the illuminant gives rise to a particular disparity value. To accomplish this, a sequence of grid patterns are projected. With each successive iteration the projection pattern is divided into finer grids. Each iteration consists of four frames. In each frame exactly one quarter of the screen is illuminated. After each iteration all parts

of the screen have been illuminated exactly once. After an iteration each of the illuminated areas are again divided into four parts and each part is illuminated in a different frame in the next iteration. The process terminates when the resolution limit of the screen and the Triclops camera is reached. In our case, this occurs after four iterations. The four iterations are shown in Figure 5. Based on the

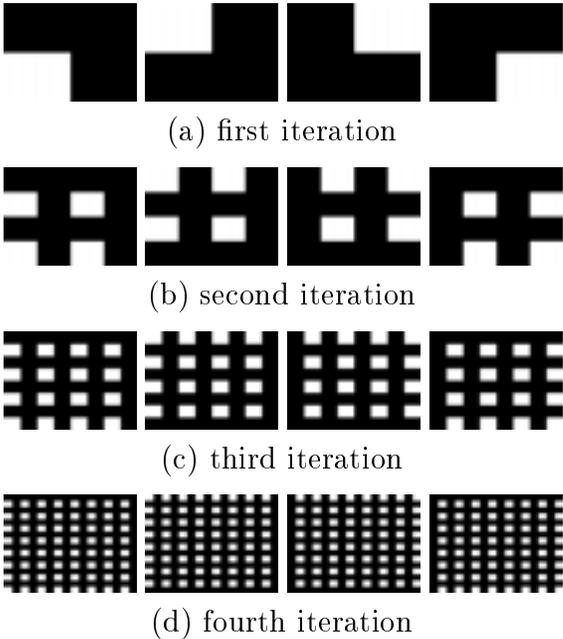


Figure 5: Coded light patterns used to identify emitted light.

images obtained by Triclops, we are now able to determine, for each pixel in the image, the region on the computer screen which emitted the illuminant. After each iteration we can determine which frame illuminated each pixel by choosing the frame in which the pixel’s intensity is brightest. This method enables us to constrain the portion of the screen from which the light was emitted in a systematic manner (see Figure 6). Through this process the possible source of

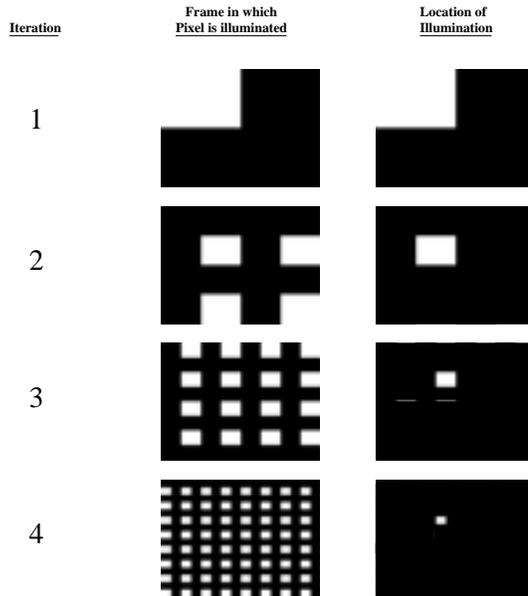


Figure 6: Example illustrating how the location of illumination is determined.

the illumination is reduced to one illuminant square in the final iteration. The center of this surviving region is used as an approximation of a point illuminant. The error in this assumption depends on the size of the regions illuminated in the final iteration as well as the distance of the reflected light from the emitted light. A derivation of the error in this approximation can be found in [33].

5.5 Dealing with Non-Planar Surfaces

The passive target technique assumes that the target is a single, infinitely large, planar surface. In reality, specular surfaces are neither infinitely large nor perfectly planar. Here we consider the issue of nonplanarity. Instead of assuming that all recovered surface points belong to a single plane, the points are divided into local regions and each region is assumed to be locally planar. To simplify

the description, we first explore the recovery of surfaces that are purely specular, with no diffuse components.

Since all the image points are specular, the points on the object’s surface can be recovered using the method described in the preceding subsections. Once the points are recovered, a set of overlapping regions are constructed and a depth value is obtained for each region by averaging depths for that region. The i^{th} region, r_i , is defined by its four bounding planes. Assuming the camera faces in the $-z$ direction and the image plane lies parallel to the xy-plane,

$$r_i = \{x_{i_{left}}, x_{i_{right}}, y_{i_{bottom}}, y_{i_{top}}\}$$

The location of region r_{i+1} , is calculated as

$$x_{i+1_{right}} = x_{i_{right}} + size_x * (1 - overlap)$$

$$x_{i+1_{left}} = x_{i+1_{right}} + size_x$$

$$y_{i+1_{bottom}} = y_{i_{bottom}} + size_y * (1 - overlap)$$

$$y_{i+1_{top}} = y_{i+1_{bottom}} + size_y$$

where $size_x$ and $size_y$ represent the size of regions measured along the x and y axis respectively. These are obtained by dividing the recovered points into n regions along each axis. To avoid allowing outlying points to cause $size_x$ or $size_y$

to be very large, we require only 95% of the recovered points to be encompassed by the area selected. The variable overlap is the percentage of the region that overlaps previous regions along each axis. We achieved reasonably good results using $n=6$ and $\text{overlap}=0.25$. The points in a region are output only if the number of points is greater than a prescribed minimum (10). Clearly, this method works most effectively for fronto-parallel surfaces.

Now consider surfaces with both specular and diffuse regions. Since we are modeling non-planar surfaces, the region of interest is defined less clearly. For example, if the object has two planes at an angle to one another, the projected pattern could be reflected in each plane. We will refer to the object we intend to model as Θ . In non-planar surfaces, diffuse points identified by the system can arise in several ways.

1. The reflection in Θ of objects located behind the projector.
2. The existence of diffuse areas on the surface of Θ .
3. A diffuse object lying between the projector and Θ .

We wish to model the diffuse points that arise due to (2) and discard those points that arise due to (1) and (3).

In order to reconstruct a specular surface with diffuse components the method is applied in two passes. In the first pass the algorithm is run exactly as described (i.e. only points initially identified as specular are used). The algorithm is then run a second time. This time the points identified as diffuse by the system are

included in solving for the surface. For each region i there are two possibilities after running pass one of the algorithm.

The first pass defined a plane in region i In this case, any diffuse points significantly behind or in front of the plane defined by the specular points are discarded. Any other diffuse points are assumed to be on the surface of the object. The current region is recalculated incorporating all of the diffuse points and all the specular points belonging to the region. Robust outlier detection is then applied to the region where the points are assumed to belong to a locally planar region.

The first pass did not define a plane in region i In this case, we look at the closest neighbor to the current region whose plane was defined. We then continue as described above, using this nearby plane to determine which of the diffuse points to include or exclude in the calculation of a plane for the region.

6 Experimental Results

6.1 Experiment 1: Planar Specular Surface

Figures 7 and 8 provide detailed results for the recovery of an ideal planar specular surface. Figure 7(a) shows the specular surface being recovered. Only the illuminated portion of the surface that is shown outlined by a red rectangle is

actually reconstructed. Figure 7(b), (c) and (d) show front, top and side views of the actual plane (in black) along with the recovered points (in red). Figure 7(d) also shows the reflected points (in blue). Figure 7(e) shows a front view of the reconstructed surface broken down by the perpendicular distance of the points to the actual plane. Similarly, Figure 8 plots the number of points recovered as a function of their perpendicular distance (error) from the actual plane.

6.2 Experiment 2: Planar Specular Surface - Multiple Orientations

Several experiments have been conducted to examine the performance of the algorithm on different planar specular surfaces.

Yaw We positioned a planar specular object, mounted on optical bench hardware, and rotated it at several different angles θ about the vertical. We then recovered points on the surface of the object and fit them to a plane. The angle of the recovered plane from the camera normal was measured and compared to the angle of rotation. The error in the orientation of the recovered plane is shown in the graph in Figure 9.

Pitch Similarly, a planar specular surface was positioned at different angles α about the horizontal. As before, points were recovered and fit to a plane. The tilt of the recovered surface from the plane of the camera was compared to the actual angle of the surface. The error in the orientation of the recovered plane is

shown in the graph in Figure 10.

6.3 Experiment 3: Planar Specular and Diffuse Surface

Figure 12(a) shows an ideal planar mirror with a diffuse target (the word “Hi”). Figure 12(b), (c) and (d), show front, top and side views of the surface respectively with the specular measurements shown in red and the diffuse measurements shown in green. The black plane seen in (c) and (d) is the actual surface and the blue points shown in (d) are the points recovered by the Triclops system. The final image, (e) shows a front view of the reconstructed surface broken down by the perpendicular distance of the points to the actual plane. Figure 11 shows

Table 1: Statistics on recovered planar surface with diffuse components.

| | |
|---------------------------------|--------------------------------|
| Total points | 2592 |
| Diffuse points | 221 |
| Specular points | 2371 |
| Mean distance from actual plane | 0.0056 <i>m</i> |
| Variance | 0.000015 <i>m</i> ² |
| Standard deviation | 0.0039 <i>m</i> |

the number of points recovered as a function of their perpendicular distance (error) from the actual plane. The distribution of the error is very similar for the specular and diffuse components (see Table 1).

Table 2: Statistics on recovered multiple planar surfaces at an angle.

| | Plane 1 | Plane 2 | Both planes |
|---------------------------------|--------------------------------|--------------------------------|--------------------------------|
| Total points | 1816 | 3141 | 4957 |
| Diffuse points | 112 | 0 | 112 |
| Specular points | 1704 | 3141 | 4957 |
| Mean distance from actual plane | 0.0074 <i>m</i> | 0.0089 <i>m</i> | .0084 <i>m</i> |
| Variance | 0.000031 <i>m</i> ² | 0.000054 <i>m</i> ² | 0.000046 <i>m</i> ² |
| Standard deviation | 0.0056 <i>m</i> | 0.0073 <i>m</i> | 0.0068 <i>m</i> |

6.4 Experiment 4: Multiple Planar Surfaces

This experiment has two planar surfaces at approximately a 30° angle to one another. The results are shown in Figure 13. Figure 13(a) shows the actual surfaces. The red rectangle outlines the area of interest as defined by the system. Figure 13(b), (c) and (d) show front top and side views of the surface with the recovered specular surface shown in red and the recovered diffuse surface shown in green. The black plane seen in (c) and (d) is the actual surface and the blue points shown in (d) are the points recovered by the Triclops system. The final image, (e) in Figure 13 shows a front view of the reconstructed surface broken down by the perpendicular distance of the points to the actual plane. Table 2 provides some statistics on the reconstruction broken down by the two planes labeled in Figure 13(a).

7 Discussion

This paper describes a novel approach to reconstructing the surface structure of large, smooth objects with highly specular surfaces. The intended primary application for the system is the *in-situ* reconstruction of man-made orbital objects. Many other applications are possible, however, it was with this space application in mind that the system was constructed.

The Coded Target Specular Stereo algorithm produces range data for highly specular surfaces. This algorithm accomplishes this by using stereo (trinocular) vision in combination with a structured light technique. The algorithm relies on the geometry of planar mirrors to transform these virtual points into points on the surface of the object. Results obtained demonstrate that the system can produce local reconstructions of simple, highly specular objects. The prototype system was able to produce range data for planar and multiple planar surfaces. The surfaces explored included those that were completely specular and those that had diffuse components. The points recovered were generally accurate to within an average of ± 0.005 m of the actual surface.

7.1 Future Work

Various improvements could be made to the system. Immediate improvements could be achieved by using a more sophisticated camera system to achieve clearer images with better resolution. There are many different directions that could be

taken in order to continue this research. A few possibilities are presented below:

- A second generation system could be built. This system could have the cameras built directly into the projection screen. The benefits of this would be twofold. Firstly, the projection screen would be directly in front of the cameras instead of off to one side; and secondly, the camera-screen calibration would only have to be done once, since the camera-screen configuration would be fixed. Even though the camera was attached to the illuminant in our system, this attachment “changed” on occasion resulting in the need to recalibrate. The system could also be equipped with four cameras instead of three. This would increase the accuracy of the correspondence task if Kanade’s multi-baseline method is used. Finally, the next system should be portable and lightweight. It should be built in a manner that will enable it to be easily mounted on a robotic arm to allow free movement of the system. This would enable the system to reconstruct complete three-dimensional representations of objects.
- In reality, most surfaces are neither perfectly specular nor perfectly diffuse but a hybrid of these two surface types. The method explored in this thesis is not designed for surfaces of hybrid reflectance. In order to adapt the algorithm to this type of problem a surface would have to be divided into three surface types; namely, specular, diffuse and hybrid. Different methods could be applied to recover points from each surface type.

- In this paper the reconstructed surfaces are presented simply as an unorganized set of coordinate points. The next generation system could use graphics methods to present the object in a more complete manner. In addition, if the system is held in a robotic arm and moved, work will have to be done in consolidating a sequence of local reconstructions.
- Further testing should be performed, preferably on orbit, on the reflective material used to cover man-made orbital objects. Man-made satellites have additional features besides specularly that complicate *in-situ* inspection. For example, many satellites spin, typically at a rate of fifty to one hundred revolutions a minute, in order to maintain a stable orbital position [28]. This motion would complicate the inspection process particularly if the system is held by a stationary robotic arm. In addition, localization of the orbital object would be a problem. Embedding the system into a robot that maneuvers along the surface of the orbital object could help to alleviate these difficulties. Such technology is being developed by researchers to assist with the inspection of airplane skins [40]. Finally, external factors, such as glare from the sun have not been considered.

8 Acknowledgments

We would like to thank the Natural Sciences and Engineering Research Council of Canada (NSERC), The Canadian Space Agency, CRESTech and MD Robotics

for their generous financial support.

References

- [1] A. Blake and G. Brelstaff. Geometry from specularities. In *Second International Conference on Computer Vision*, pages 394–403, 1988.
- [2] E. Boyer and M. O. Berger. 3d surface reconstruction using occluding contours. *International Journal of Computer Vision*, 22(3):219–233, March/April 1997.
- [3] T. Bulow. Consistent phase estimation for recovering depth from moiré interferograms. Technical Report MS-CIS-01-14, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, 1991.
- [4] S.-Y. Cho and T. W. S. Chow. Neural computation approach for developing a 3-d shape reconstruction model. *IEEE Transactions on Neural Networks*, 12(5), September 2001.
- [5] M. Clerc and S. Mallat. The texture gradient equation for recovering shape from texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):536–549, 2002.
- [6] L.S. Davis, L. Janos, and S. M. Dunn. Efficient recovery of shape from texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(5):485–492, 1983.

- [7] D. Halliday, R. Resnick, and J. Walker. *Fundamentals of Physics Extended Fifth Edition*. John Wiley & Sons, Inc., New York, USA, 1997.
- [8] M. A. Halstead, B. A. Barsky, S. A. Klein, and R. B. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *SIGGRAPH '96*, New Orleans, Louisiana, August 1996.
- [9] S. Hata. Shape detection of small specular surface using color stripe lighting. In *International Conference on Pattern Recognition*, pages 554–557, 1992.
- [10] B. K. P. Horn. *Shape from Shading: A method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD thesis, Massachusetts Institute of Technology, Boston, Massachusetts, USA, 1970.
- [11] B. K. P. Horn. *Robot Vision*. MIT Press, Cambridge, Massachusetts, USA, 1986.
- [12] K. Ikeuchi. Determining surface orientations of specular surfaces using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6):661–669, 1981.
- [13] Y. Iwahori, R. J. Woodham, H. Tanaka, and N. Ishii. Neural network to reconstruct specular surface shape from its three shading images. In *Proceedings of the International Joint Conference on Neural Networks*, pages 1181 – 1184, 1993.

- [14] M. Jenkin and P. Jasiobedzki. Computation of stereo disparity for space materials. In *IEEE/RSJ IROS*, Victoria, Canada, July 1998.
- [15] T. Kanade, H. Kato, S. Kimura, A. Yoshida, and K. Oda. Development of a video-rate stereo machine. In *IEEE/RSJ IROS*, Pittsburgh, USA, August 1995.
- [16] T. Kanade and M. Okutomi. A stereo matching algorithm with an adaptive window: Theory and experiment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(9):920–932, September 1994.
- [17] S. Kiyasu, H. Hoshino, K. Yano, and S. Fujimura. Measurement of the 3-d shape of specular polyhedrons using an m-array coded light source. *Transactions on Instrumentation and Measurement*, 44(3):775–778, June 1995.
- [18] R. Klette, A. Koschan, K. Schluns, and V. Rodehorst. Evaluation of surface reconstruction methods. In *Proceedings of New Zealand Image and Vision Computing*, pages 3 – 12, Lincoln, Canterbury, August 1995.
- [19] R. Klette, K. Schluns, and A. Koschan. *Computer Vision Three-Dimensional Data from Images*. Springer-Verlag Singapore Pte. Ltd., Singapore, 1998.
- [20] K. Koshikawa. A polarimetric approach to shape understanding of glossy objects. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 493–495, 1979.

- [21] J. Lu and J. Little. Reflectance and shape from a rotating object. Technical Report 95-9, Department of Computer Science, University of British Columbia, Vancouver, BC, April 1995.
- [22] C. Marinos and A. Blake. Shape from texture: The homogeneity hypothesis. In *Third International Conference on Computer Vision*, pages 350–353, 1990.
- [23] M. Maruyama and S. Abe. Range sensing by projecting multiple slits with random cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(6):647–651, 1993.
- [24] McGraw-Hill, USA. *McGraw-Hill Dictionary of Scientific and Technical Terms*, fifth edition, 1994.
- [25] S. K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation*, 6(4), August 1990.
- [26] S.K. Nayar, X.-S. Fang, and T. Boult. Separation of reflective components using color and polarization. Technical Report CUCS-058-92, Department of Computer Science, Columbia University, New York, New York, 1993.
- [27] J.-S.. Park and J. T. Tou. Highlight separation and surface orientation for 3-d specular objects. In *Proceeding of the International Conference on Pattern Recognition*, pages 331–335, 1990.

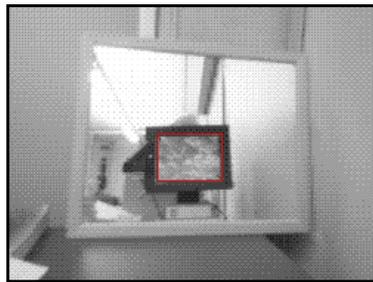
- [28] B. Pattan. *Satellite Systems Principles and Technologies*. Van Nostrand Reinhold, New York, USA, 1993.
- [29] Point grey research inc. <http://www.ptgrey.com>.
- [30] S. Rahmann. Polarization images: A geometric interpretation for shape analysis. In *Proceedings of the 15th International Conference on Pattern Recognition*, volume 3, pages 538–542, 2000.
- [31] S. Rahmann and N. Canterakis. Reconstruction of specular surfaces using polarization imaging. In *Proc. Conf. on Computer Vision and Pattern Recognition*, pages 149–155, 2001.
- [32] K. V. Rajaram, G. Parthasarathy, and M. A. Faruqi. A neural network approach to photometric stereo inversion of real-world reflectance maps for extracting 3-d shapes of objects. *IEEE Transactions on Systems, Man, and Cybernetics*, 25(9):1289–1300, September 1995.
- [33] A. Ripsman. Local surface reconstruction of orbital objects. Master’s thesis, York University, Toronto, Canada, 2002.
- [34] H. Saito, K. Omata, and S. Ozawa. Recovery of shape and surface reflectance of specular object from rotation of light source. In *Proceeding of the Second International Conference on 3-D Digital Imaging and Modeling*, pages 526–535, 1999.

- [35] J. Salvi. *An Approach to Coded Structured Light to Obtain Three Dimensional Information*. PhD thesis, University of Girona, Girona, Spain, 1999.
- [36] 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):44–55, January 1988.
- [37] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. Technical Report CMU-CS-92-207, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, November 1992.
- [38] S. Savarese and P. Perona. Local analysis for 3d reconstruction of specular surfaces. In *Proc. Conf. on Computer Vision and Pattern Recognition*, volume 2, pages 738–745, 2001.
- [39] S. A. Shafer and T. Kanade. Using shadows in finding surface orientations. In L. B. Wolff, S. A. Shafer, and G. E. Healey, editors, *Shape Recovery*, pages 346–377. Jones and Bartlett Publishers, 1992.
- [40] M. Siegel. Remote and automated inspection: status and prospects. In *First Joint DoD/FAA/NASA Conference on Aging Aircraft*, Ogden, UT, USA, July 1997.
- [41] F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. In *Proc. Conf. on Computer Vision and Pattern Recognition*, Champaign, Illinois, June 1992.

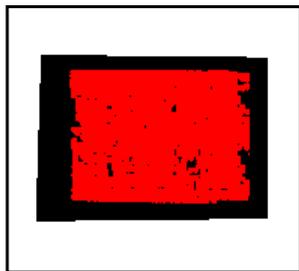
- [42] M. Tarini, H. Lensch, M. Goesele, and H.-P. Seidel. Shape from distortion: 3d range scanning of mirroring objects. In *SIGGRAPH*, S. Antonio, Texas, USA, 2002.
- [43] Y.-L. Tian and H. T. Tsui. Shape from shading for non-lambertian surfaces for one color image. In *International Conference on Pattern Recognition*, pages 258–262, 1996.
- [44] J. Tolson. <http://www.sfsu.edu/~jtolson/textures/sor/pink/pinkmix/4.htm>.
- [45] P. Vuylsteke and A. Oosterlinck. Range image acquisition with a single binary-encoded light pattern. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):148–164, 1990.
- [46] L. B. Wolff. Spectral and polarization stereo methods using a single light source. In *International Conference on Computer Vision*, pages 708–715, London, England, 1987.
- [47] L. B. Wolff. Surface orientation from two camera stereo with polarizers. In *Proc. Opt. Illumination Image Sensing Machine Vision IV*, volume 1194, pages 287–297, Philadelphia, PA., November 1989.
- [48] L. B. Wolff and T.E. Boulton. Constraining object features using a polarization reflection model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):635–657, 1991.

- [49] R. J. Woodham. Photometric stereo: A reflectance map technique for determining surface orientation from a single view. In *Proceedings of the SPIE 22nd Annual Technical Symposium*, San Diego, California, USA, August 1978.
- [50] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [51] C. Wust and D. W. Capson. Surface profile measurement using color fringe projection. *International Journal on Machine Vision and Applications*, 4:193–203, 1991.
- [52] Y. Xiong and S. A. Shafer. Depth from focusing and defocusing. Technical Report CMU-RI-TR-93-07, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, March 1993.
- [53] J. Y. Zheng, Y. Fukagawa, and N. Abe. 3d surface estimation and model construction from specular motion in image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(5):513–520, May 1997.
- [54] J. Y. Zheng and A. Murata. 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):913–920, August 2000.

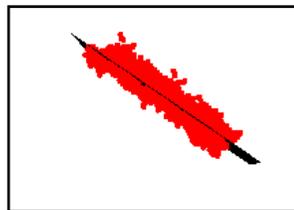
- [55] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, 7(1):38–42, 1989.



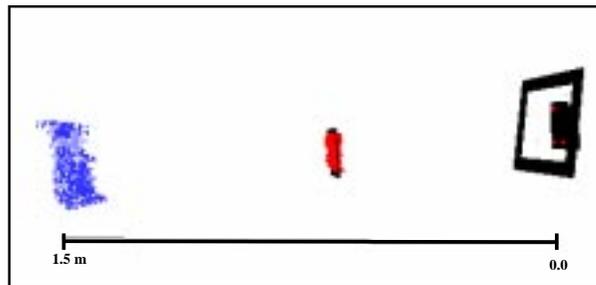
(a) actual surface



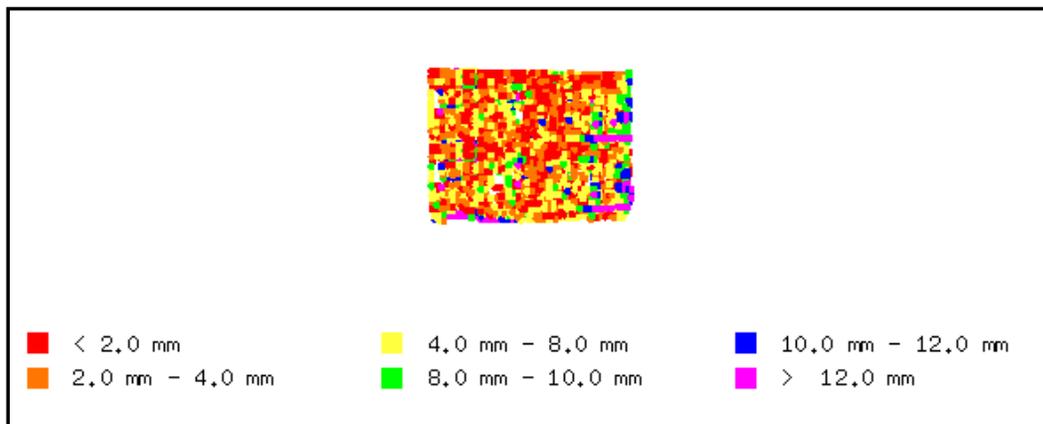
(b) front view



(c) top view



(d) side view



(e) distance from the plane

Figure 7: Reconstruction of a planar specular surface.

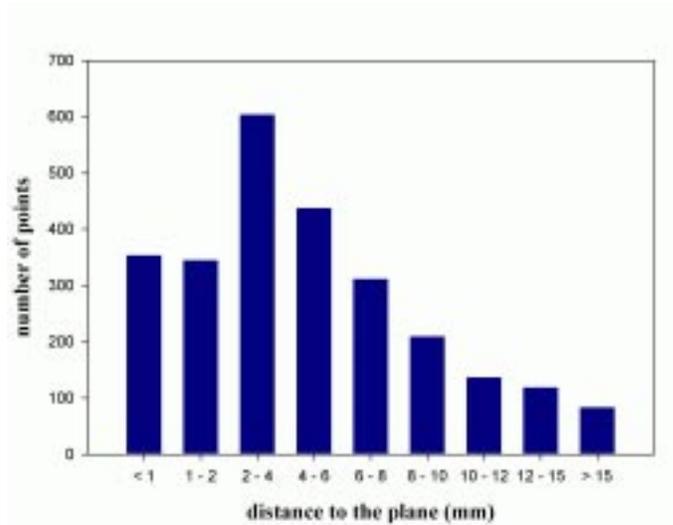


Figure 8: Histogram showing the distribution of perpendicular distance (error) between recovered specular points and the actual plane for a planar surface of approximately 11 cm by 14 cm, approximately 65 cm from the sensor. The number of points is 2594, mean distance to the plane is 0.005 m and standard deviation is 0.0042 m.

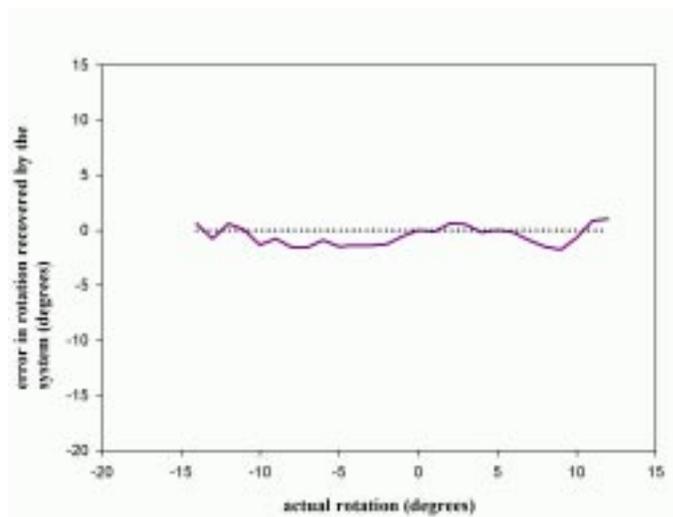


Figure 9: Error in angle of rotation, θ , recovered by the system.

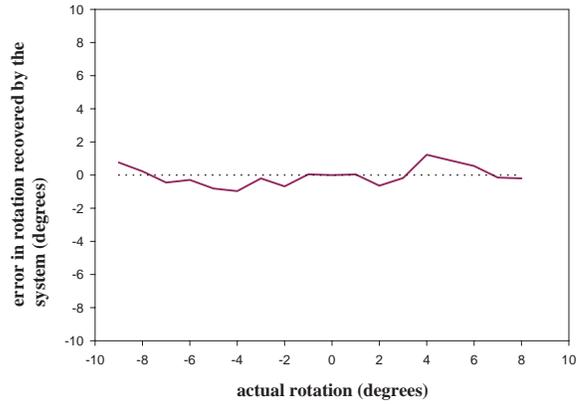


Figure 10: Error in angle of rotation, α , recovered by the system.

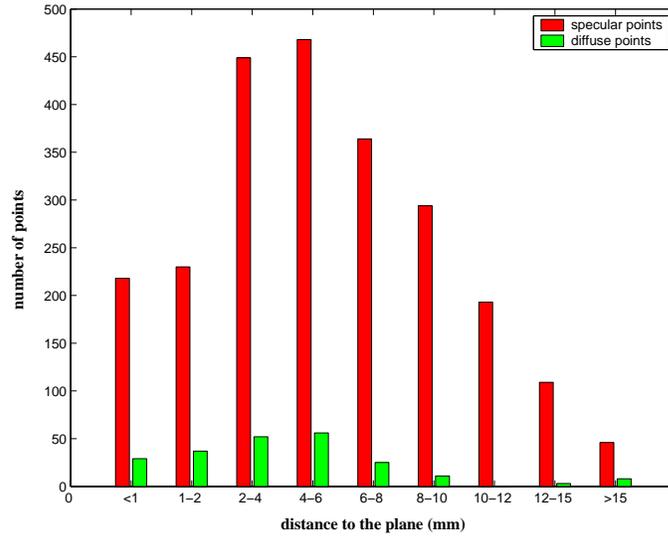


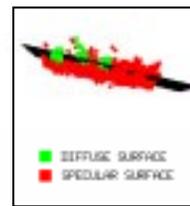
Figure 11: Histogram showing the distribution of perpendicular distance (error) between recovered specular and diffuse points and the actual plane for a planar surface of approximately 11 cm by 15 cm, approximately 50 cm from the sensor.



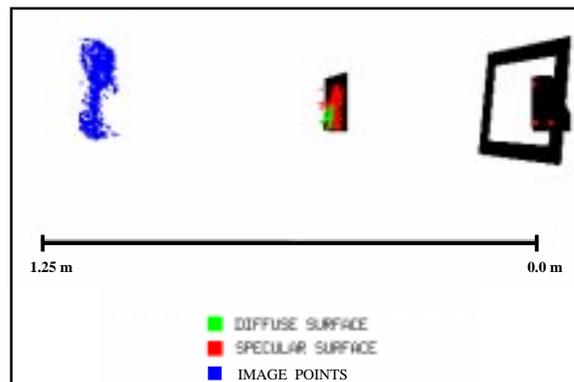
(a) actual surface



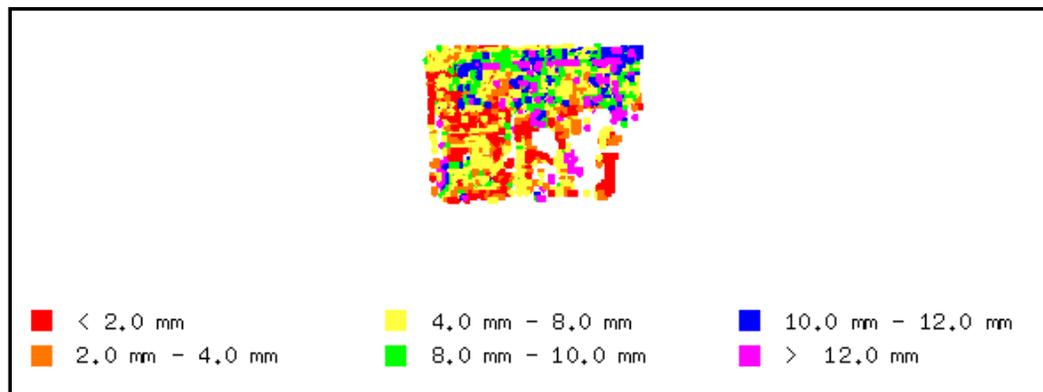
(b) front view



(c) top view

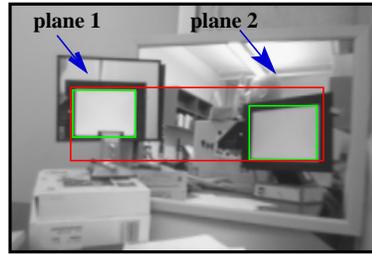


(d) side view

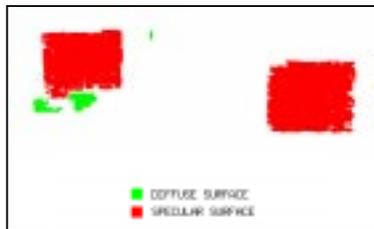


(e) distance from the plane

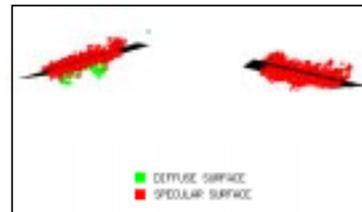
Figure 12: Reconstruction of a planar specular surface with diffuse components.



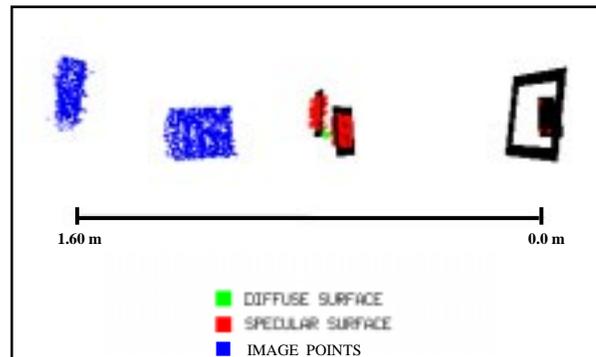
(a) actual surface



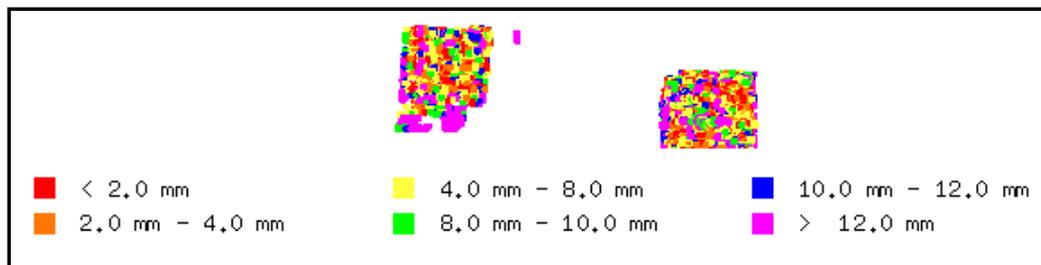
(b) front view



(c) top view



(d) side view



(e) distance from the plane

Figure 13: Reconstruction of two planar specular surfaces at an angle.