Computation of Stereo Disparity for Space Materials

M. Jenkin[1] and P. Jasiobedzki[2]

[1]York University, 4700 Keele St., North York, Ontario, Canada

[2]SPAR Aerospace, 9445 Airport Rd., Brampton, Ontario, Canada

Abstract

One of the challenges facing computer vision systems used in space is the presence of specular surfaces. Such surfaces lead to several adverse effects when imaged by vision systems such as the creation of reflected "virtual" images of objects due to specular reflections, specular reflection of light from the sun and other sources, and total reflection of any projected illuminants including laser beams. These effects may lead to incorrect measurements and loss of data in the case of sensor saturation or inadequate intensity of the returned laser beams in the case of an active illuminant. In addition, the instruments inside space structures such as satellites may be extremely sensitive to active illuminants such as laser beams or radar signals, and thus passive vision systems which rely on either natural on low-power projection systems are preferred over active sensing technologies. An additional advantage of fixed stereo based techniques over scanning rangefinders lies in the fact that fixed stereo systems do not contain moving parts that are prone to failure and expensive to qualify for space and maintain. Fixed stereo based systems use cameras, framegrabbers and computers already qualified for space but must deal with the issue of specular reflections. Here we consider the task of recovering the local surface structure of highly specular surfaces such as satellites using passive stereopsis without resulting to the introduction of additional light sources. In particular we examine the use of stereo cameras to recover surfaces which result in perfectly specular reflection.

1. Introduction

The inspection of satellites and other orbiting device becomes increasingly important as more and more material is placed in orbit. In order to maximize the longevity of such devices, to provide for *in situ* upgrades, etc., it is necessary to periodically inspect these devices for damage and to perform servicing tasks such as the exchange of Orbital Replaceable Units (ORUs). Access to ORUs requires approaching the spacecraft, locating and opening protective enclosures, and grappling and replacing modules. Accomplishing these tasks involves surveying the structure for damaage and identifying access panels. As was made clear on a recent mission to upgrade the Hubble telescope, not all damage is easy to detect at a distance, and sophisticated sensing mechanisms and sensors are required in order to survey/inspect space structures before repairs can be carried out. As satellites are relatively fragile devices and as they tend to be sensitive to various forms of radiation, any such inspection should be carried using either natural illumination or using illumination which is benign to onboard sensors, heat shields, etc.

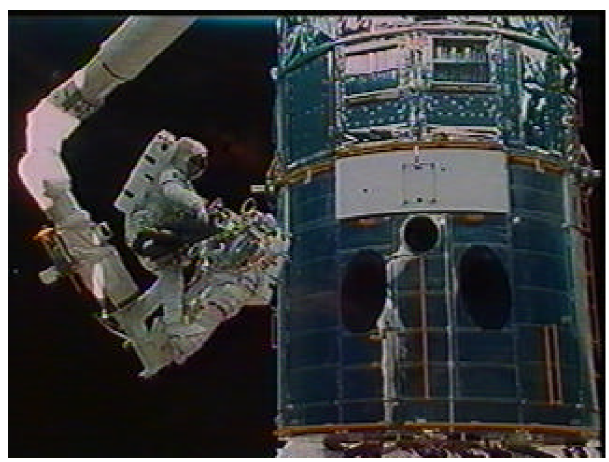

Figure 1. Typical space scene. Here an astronaut maneuvers on the end of the Canada Arm to service the Hubble Space telescope. Note the highly reflective surface on the telescope (the astronaut is reflected in it.)

An additional constraint on the type of sensor that can be deployed is related to the cost associated with the use of sensors with moving parts. Such devices are prone to failure and expensive to

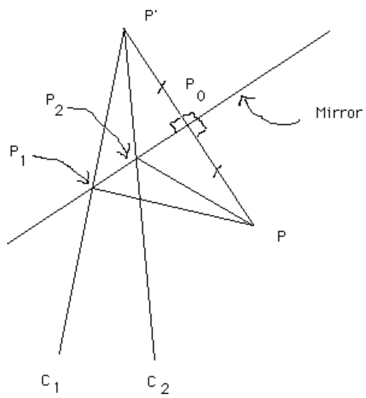

Figure 2: Plane mirror

maintain and qualify for space. Systems which use standard cameras, framegrabbers, and computers which are already space qualified are to be preferred. Given all of these constraints, systems such as passive stereopsis, which rely on either natural illumination or immobile low-power illuminant-based approaches, are to be preferred over active, high-power illuminant techniques.

The inspection of an orbiting device should recover metric information concerning the shape and structure of the device so that the measured structure can be compared against "as delivered" blueprints of the device in order to check for warping/expansion of the entire structure, as well as identifying local changes in surface structure. Any visual inspection of a satellite is complicated by the composition of the external coverings of the vehicle. Space vehicles are typically covered in either thermal blankets, made of relatively uniform and featureless material, or in highly reflective material used for thermal reflecting and radiating surfaces and solar arrays (see Figure 1). The reflective materials promote specular reflections and thus local surface intensities do not necessarily correspond to local surface material properties but rather correspond to the intensity of the illuminant. This is clearly evident in Figure 1 in which the astronaut is reflected in the surface of the Hubble telescope. Any mechanism for visual inspection of such surfaces will have to deal with the reality that recovered image structure may only be related indirectly to true surface structures in the scene.

The classic passive technique for recovering metric 3-D object structure is based on stereopsis. See Marr and Poggio (1979), Grimson (1981), Howard and Rogers (1995) and others for some classic approaches and a review of the literature. Computing stereo disparity involves determining the correspondence between the projection in a pair of cameras of a single entity in three dimensional space. Given this correspondence, and the known camera geometry, it is possible to recover the location of the three dimensional entity through triangulation. Correspondence is known to be the hard problem in stereopsis and many different approaches have been suggested. Classical computational approaches to the task of determining stereo correspondence are to either (a) detect features in each image and then to match the features, or (b) to match local brightness distributions between the images. Underlying these approaches to the correspondence problem is the assumption that a scene feature gives rise to identical image brightness patterns in the two cameras which via triangulation correspond to true surfaces in the scene. Unfortunately, given the highly reflective properties of space materials, it is unlikely that this assumption will hold in practice in outer space, resulting in false disparity values and incorrect depth estimates if classic stereopsis algorithms are applied blindly to highly specular surfaces.

Surfaces of space structures will give rise to both specular and diffuse reflections. Here we conisder only the task of determining the local surface structure of purely specular surfaces.

2. Specular Surfaces

Specular surfaces pose a number of unique problems for stereopsis algorithms. The primary problem is that should a match be found, it does not necessarily correspond to a point on the surface of an object. This is depicted in Figure 2, which shows the reflection of a point P in a plane mirror. If two cameras with nodal points at C_1 and C_2 view a point P in a mirror, they "view" or reconstruct this point at P'. The recovered point P' does not lie on the surface of the mirror, nor does it correspond to the true image point, rather it corresponds to the “virtual” reflection of P in the mirror.

Now suppose that the surface has both specular and diffuse components. Then we should expect some binocular matches to correspond to features which lie on the surface of the object, while other matches may lie beneath the surface (convex mirrors) while others may lie above the surface (concave mirrors).

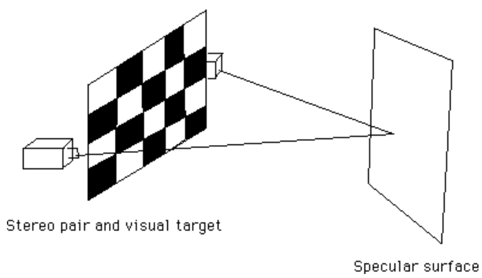

Figure 3. Passive visual target. The visual target located with the cameras is designed to be easily detected using whatever stereo algorithm is being used.

Although recovery of P' does not provide direct information concerning the surface, it can be used to produce surface information indirectly under the assumption that the mirror is locally planar between P_1 and P_2. If P is known in addition to P', then we can determine P_0 by noting that $P_0 = (P + P')/2$. As PP' is perpendicular to the planer mirror segment that lies betwen P_1 and P_2, we can use the rays C_1P' and C_2P' to recover P_1 and P_2. Note that this works even if the mirror is not completely planar - it only needs to be planar in the region near P_1P_2.

A true point P and the reconstruction of its reflection P' provides an indirect method for the recovery of highly specular surfaces. This suggests two potential approaches for obtaining surface information concerning specular surfaces;

1. Reflect a known target in the specular surface for which P is known *a priori*, and then determine P' via classic stereo mechanisms or by constructing the target in such a way that stereopsis is trivial, e.g., by modulating the intensity of the known target.
2. If an object *and* its reflection are both visible in the scene, then it may be possible to solve the stereo problem for the object and for the reflected version of the object and then solving for the correspondence between the object and its reflection.

The following sections provide some preliminary results obtained under each of these two approaches

3. Reflection of a known target

A special target was constructed to be reflected in the specular surface. The concept is sketched in Figure 3, while Figure 4 show the experimental setup. Two convergent CCD cameras are mounted with a patterned surface between them. In the prototype system described here a grid pattern of points is used, but patterns which would be trivially recovered in an automatic manner would be better suited for space deployment. In Figure 4, the cameras face a calibration object and mirror. During an initial calibration phase, the mirror is removed and the stereo rig is calibrated using a linear least-squares calibration process (Horn, 1986). This calibration process is used to define a global 3D coordinate system as well as to calibrate the individual cameras within this global coordinate system. The optical centres of the two cameras

Figure 4: Experimental setup. The stereo camera pair (on the left) views a mirror and calibration target (on the right). A visual target is mounted between the two video cameras.

are identified in the global coordinate system, as well as the coordinates of readily identified feature points on the target mounted between the two cameras. A specular test object, here a mirror, is then placed in the scene (see Figure 4).

Figure 5 shows the binocular view of the target. A highly specular surface (here a mirror) is viewed by the camera pair. The passive target placed between the two cameras is clearly reflected in the highly specular surface of the mirror. Correspondences between the reflected points of the target in the two images are constructed. (In the test here these correspondences have been performed by hand. In a production system a token matching algorithm would have to be implemented). These correspondences, coupled with the known calibration, allows the computation of points on the surface of the mirror. As shown in Figure 2 and as discussed earlier, for each visible point on the target, we can solve for the intersection of the plane defined by P and P' and the lines defined by the optical centres of the two cameras and the reflected image point. Each feature point on the target gives rise to two points on the specular surface.

Figure 6 shows the results of the system in operation. It renders the thre dimensional environment within which the system operates as well as the the recovered image points. The 'L'-shaped object in the center of the scene is the calibration object. This object defines the global coordinate system within which the system operates. Straight lines connect the nodal points of the cameras and the fixation point. As can be seen from the image the cameras fixate a point slightly in front of the fixation point. Points on the visual target lie between the nodal points of the two points and are plotted. The reflection of

Figure 5: Stereo pair of a mirror. The reflection of the target mounted between the two cameras is clearly visible.

theses points (their mirrored image) which are the result of the stereopsis correspondence process are plotted. These appear behind both the mirror as well as the calibration object. Clearly these points do not correspond to the surface of the mirror which lies between the cameras and the calibration object. Points "on" the mirror are inferred given a point on the visual target and its recovered mirror image. These points completely obliterate the mirror which was drawn in Figure 6 at its appropriate location.

An examination of the mirror points identifies two outliers which are a result of inaccuracies in the hand solution of the correspondence problem. In fact they correspond to a single poor correspondence which is visible behind the calibration object.

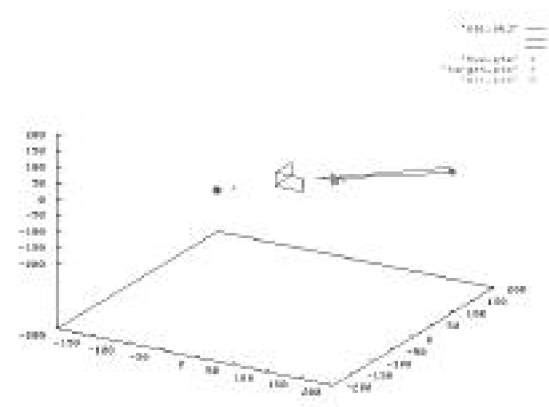

Figure 6: 3d envrionment showing cameras, the mirror, real points and their reflection. Points which have been hypothesized to lie on the mirror are also shown.

The general technique of reflecting a special-purpose target in a specular suface to recover the shape of the surface indirectly can be quite powerful. Special targets can be designed to completely eliminate the correspondence problem -- the target could consist of individually accessable LED's, for example -- and each feature point obtains two points on the specular surface.

Of course the technique's weakest facet is the fact that it requires a special purpose target to be

Figure 7. Simple scene containing a mirror

constructed and calibrated. An alternative would be to use objects which exist in the scene to infer the structure of specular surfaces. This is considered in the following section.

4. Reflection of an arbitrary target

The technique described above requires a known target in order to recover the specular surface An alternative is to use unknown arbitrary objects in the scene itself in order to recover the mirrored surface. . To examine the potential of this type approach a scene was constructed with a number of objects and a large mirror (see Figure 7). Note that Figure 7 does not contain two cerial boxes, one is a reflection of the other.

Figure 8 shows the reconstructed 3D model obtained when Figure 7 is processed to recover disparities without modelling reflective surfaces. As would be expected, the reconstructed scene contains two boxes and two spaceships on a continuous planar surface of the table. As there are a number of occlusions in the scene some parts of the table are missing in the reconstruction.

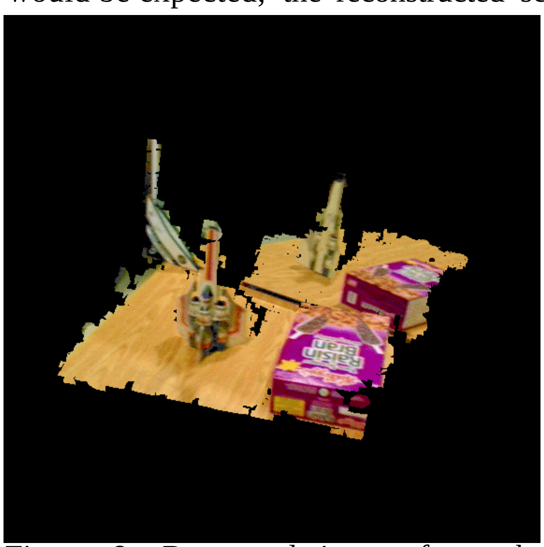

Figure 8. Recovered image from the stereo pair in Figure 7. To recover the mirror, matches between objects and their reflections must be consructed.

Two main effects can be seen in this reconstruction:

1. A number of objects appear twice in the scene. For example, the top surface of the box appears twice in the reconstruction (the real box and it's mirror reflection). In this case, the same point appears twice in each image of the stereo pair (once as real point and once as a virtual one). This suggests a potential approach for recovering specular surfaces in a scene: first process the imagery with a classic stereopsis algorithm searching for binocular matches, then examine the set of binocular matches for correspondences within the scene. This second matching stage searches for objects which appear twice -- once for the real object and once reflected. Note that this correspondence process has much of the same flavour as binocular matching but it is a three dimensional problem (rather than two), and the global matching stage is somewhat more complex as there may exist many possible matches. Imagine a picket fence reflected in a mirror, for example.

2. Not everything visible in the mirror is visible in the real image and vice-versa. For example, in the reconstruction of the real image the top side of the spaceship is visible, while in the reconstruction of the mirror images the bottom side is visible. In this case it is not possible to recover information about the mirror as the same material does not appear in both real and virtual form. Although this connection between the top and bottom side of the spaceship in the 3D reconstruction cannot be made without additional information, it would certainly be possible to perform this connection if additional information was availble. For example, this connection could be accomplished if an *a priori* model of the object was availble. Then the specular surface may be recovered by matching both the real and virtual objects in the scene. Following a similar geometric consideration as before but at a level of objects instead of individual points, the geometry of the situation is described in Figure 9.

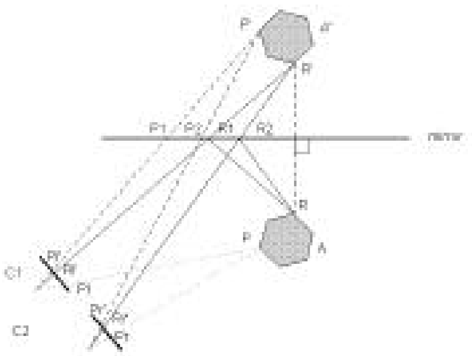

Figure 9: Model matching

It is not possible to establish direct correspondence between the points R and R’ because only R’ can be seen by both cameras. Assuming a model of the observed object, A can be matched based on visible faces of the real object. Then the pose of the object A can be recovered and the location of point P can be inferred. Similarly, if the model can be matched based on the visible faces of the virtual object A’, then the pose of A’ can be recovered and the hidden portions of the virtual object can be matched with visible portions of the real object. Matching of points on A with points on its mirror reflection A’ provides constraints on the location of the specular surface. Except in certain degenerate cases (mirror symmetric models, for example), it is possible to identify which of the objects (A and A’) is real and which one is virtual. Matching of the model with the reconstructed data can be assisted by using, for example, points such as P and P’ visible in simultaneously in both images.

The remaining computational problem is, of course, the problem of finding points or recognizing objects and solving for their correspondence in the mirror.

4. Discussion

This paper has looked at two different approaches - a passive target approach and a fully passive approach -- for the recovery of locally planar specular surfaces from a passive stereo viewer. These approach take advantage of an assumption of local mirror planarity in order to recover two points on the mirror for each object-reflection match. Both approaches have been examined under laboratory conditions. The passive target approach relies on the reflection of a known target in order to infer the local structure of the specular surface. The fully passive approach does not use a known target but rather must solve the stereo correspondence problem as well as solving the real-virtual correspondence task. Although neither approach has been implemented fully, a passive target system has been constructed which relies on human selection of corresponding points, while the fully passive system relies on human selection of real-virtual correspondences.

The passive target approach is likely to be of more use in an industrial environment as a specially designed target can be used to simplify the problem of finding specular surfaces. The construction of an active target which would simplify machine selection of stereo correspondences is currently being investigated.

Bibliography

Blake, A., and Bulthoff, H., “Shape from specularities: computation and psychophysics”, *Phil. Trans. R. Soc. Lond. B*, 331: 237-252, 1991.

Faugeras, O., *Three dimensional computer vision*, MIT Press, Cambirdge, MA, 1993.

Grimson, W. E. L., “A computer implementation of a theory of human stereo vision”, *Phil. Trans. Roy. Soc.*, B292, 217-253, 1981.

Horn, B., and Schunk, B., “Determining optical flow”, *Artificial Intelligence* , 17, 185-204, 1981.

Horn, B., *Robot Vision* , MIT Press, Cambridge, MA 19 81.

Howard, I. P., and Rogers, B. J. *Binocular vision and stereopsis*, Oxford University Press, 1995.

Jenkin, M., Jepson, A., Tsotsos, J., Techniques for disparity measurement, *CVGIP:IU*, 53(1), 14-30, 1991.

Marr, D., and Poggio, T., “A computational theory of human stereo vision”, *Proc. Roy. Soc. Lond*, B204, 301-328, 1979.

Ozanian, T., “Approaches for stereo matching: A review”, *Modelling, Identification and Control,* 16(2), 65-94, 1995.

Thorpe, C., “An anlysis of interest operators for FIDO”, Workshop on computer vision: representation and control, 1984.

Witkin, A., Terzopoulos, D., Kass, M., “Signal matching through scale space”, Proc. 5th Nat. Conf. on AI, 714-719, Pittsburgh, PA, 1986.