Machine Vision for Autonomous Small Body Navigation[1]

Andrew E. Johnson, Yang Cheng and Larry H. Matthies
Jet Propulsion Laboratory, California Institute of Technology
Mail Stop 125-209, 4800 Oak Grove Drive, Pasadena, CA 91109
(818) 354-0357
{aej,ycheng,lhm}@robotics.jpl.nasa.gov

Abstract--This paper describes machine vision algorithms that enable precision guidance and hazard avoidance during small body exploration through onboard visual feature tracking and landmark recognition. These algorithms provide estimates of spacecraft relative motion and absolute position used to guide the spacecraft during autonomous landing and exploration. They also enable hazard avoidance by providing estimates of 3-D surface topography through processing of monocular image streams. This form of onboard autonomy is a critical enabling technology for multiple future missions including Comet Nucleus Sample Return, Large Asteroid Sample Return, Titan Organics Explorer and Europa Lander and Mars lander missions.

TABLE OF CONTENTS

1. INTRODUCTION
2. PROBLEM FORMULATION
3. ALGORITHMS
4. CONCLUSIONS

1. INTRODUCTION

Autonomous spacecraft systems have the potential to reduce costs while enhancing existing systems and enabling new capabilities for future deep space missions. For example, exploration of comets, asteroids and moons of outer planets will benefit tremendously from on-board systems that autonomously and accurately determine spacecraft position relative to a proximal small body. With such a system, complex trajectories can be followed safely and accurately in the dynamic small body environment. This capability will enable precision guidance to scientifically interesting targets, hazard avoidance, autonomous landing, and sample return with little or no human interaction.

Design of an autonomous navigation system should balance positional accuracy against the typical constraints in spacecraft design of power, mass, volume, and complexity. Cameras are proven spacecraft sensors; most spacecraft carry cameras for scientific imaging or optical navigation. The low cost, low power, low mass, and proven flight record of cameras also make them an attractive sensing solution for autonomous navigation. Camera images can be processed by the flight computer to estimate spacecraft motion, body absolute position and 3-D surface topography. By integrating spacecraft cameras with on-board processing, an autonomous navigation sensing system can be realized with little or no addition to typical spacecraft systems.

We are developing machine vision algorithms that take as input streams of images from a single nadir pointing camera and output estimates of spacecraft relative motion, spacecraft body absolute position and imaged surface topography. These estimates can be passed directly to the spacecraft guidance, navigation, and control system for following of safe and precise trajectories. Motion is computed from feature tracking followed by two-frame image-based motion estimation. Given motion estimates, our algorithms reconstruct the 3-D topography of the imaged terrain using efficient motion stereo techniques. This topography can be used to detect hazards; it can also be used to build a 3-D model of the imaged surface. Given a 3-D model, our algorithms determine the position of the spacecraft in a body relative frame by matching landmarks extracted from an image stream to those stored in the 3-D model.

2. PROBLEM FORMULATION

During algorithm development, we have placed an emphasis on robustness to noise and outliers, generality in terms of image surface characteristics and application domain, and algorithm efficiency. Before describing our algorithms, we would like to expand on the problems we are investigating and relate them to the state of the art in machine vision and autonomous spacecraft navigation.

Visual Position Estimation

Current missions require optical navigation for orbit determination and instrument pointing during close fly-bys of small bodies and moons of the outer planets. This is implemented by ground-based image processing to extract centroids of small reference targets like asteroids and moons. For the NEAR mission, orbit determination around asteroid Eros will use manual designation of known landmark features on the surface of the asteroid [8] Limited automation was introduced in the New Millennium DS-1

[1] 0-7803-5846-5/00/$10.00 © 2000 IEEE

mission by implementing onboard centroiding of reference asteroids for autonomous navigation in small body fly-bys [9]. Proposed missions to explore comets and asteroids will not be able to rely on such techniques, because safe, precise navigation will require accurate knowledge of complex surface topography and because the round-trip light time will not allow this to be done on the ground.

Although some degree of autonomous, onboard position estimation capability has been demonstrated, the feature tracking and landmark recognition capabilities required to enable safe small body exploration do not exist. One method for visual position estimation relies on tracking image features through a sequence of images. Image features are image pixels that have a high probability of being matched between two images taken of the small body surface from similar, but not necessarily the same, camera locations. By detecting and then tracking image features through a sequence of images, the relative motion of the spacecraft can be determined between frames [5]. This capability is useful for maintaining continuous estimates of spacecraft position, but since it does not give absolute position with respect to a body centered coordinate system, its usefulness is limited.

Another method for visual position estimation is landmark recognition. A landmark is a 3-D position on the surface of a body whose appearance is stable across moderate changes in viewing direction and illumination conditions (e.g., craters on an asteroid [6]). Landmarks are detected during 3-D modeling of the body and stored in a database. During landmark recognition, landmarks detected in an image are matched to landmarks in the database. Since the 3-D position of landmarks are known, recognizing a few landmarks in a single image is sufficient for determining the absolute position of the spacecraft relative to the body centered coordinate system. Landmark recognition is more time consuming than feature tracking, however, these two methods of position estimation are complimentary. By combining the continuous updates of relative position from feature tracking with the occasional updates of absolute position from landmark recognition, continuous estimates of spacecraft position in absolute body centered coordinates can be obtained.

Motion Stereo Vision

Stereo imaging has been studied extensively, and well-known techniques for reconstructing dense surfaces from stereo images exist [12]. Traditional stereo imaging (i.e., two or more rigidly attached cameras) cannot be applied directly to the small body exploration problem, except near to the surface, because at high altitude the camera baseline required for structure recovery is too large for typical spacecraft structures. However, using spacecraft motion estimates and stereo vision techniques, it is possible to generate dense topographic maps of a small body surface from monocular image streams. This technique, called motion stereo, has the advantage of being applicable at any altitude above the small body surface. However, it requires a more complicated algorithm than typical binocular stereo because the baseline between images is variable and must be computed from navigation sensor inputs. Motion stereo is an important component of autonomous small body exploration because it provides the 3-D structure needed for hazard detection and landmark recognition.

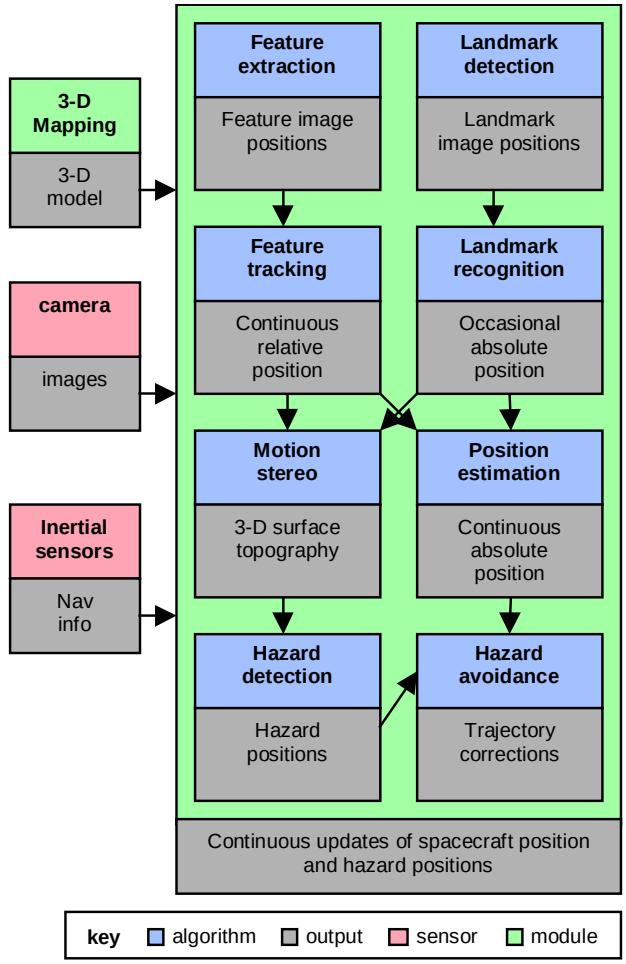

Figure 1 Algorithm block diagram.

Hazard Avoidance

Hazard detection is the process where scene topography is analyzed to detect landing sites that are unsafe for the spacecraft. Hazards can be characterized as high-level (e.g., rocks, cliffs) or low-level (e.g., local surface slope and roughness), High-level hazards are detected by segmenting hazard from the background while low-level hazards are computed at each pixel in an image. High-level hazard detection requires the definition of models for objects that promote efficient hazard detection and accurate localization of hazards. The challenge of low-level hazard detection is deciding what combinations of low-level hazards constitute a hazard for the spacecraft.

Hazard avoidance combines path planning to avoid detected hazards with constraints on fuel and spacecraft control authority to generate trajectories that guide the spacecraft to a safe landing site. Currently we are focussing on the hazard detection problem. In the future we plan to use our algorithms as a front end to a complete hazard detection and avoidance system for safe and autonomous small body landing.

3. ALGORITHMS

As shown in Figure 1 , we are developing a complete set of algorithms for passive image-based small body navigation and hazard detection. Currently we have algorithms for feature-based relative motion estimation, motion stereo vision, surface landmark-based absolute position estimation, and terrain map hazard assessment. Below we describe these algorithms and present results generated from images acquired of a comet analog in a controlled laboratory setting.

Motion Estimation

We define spacecraft motion as the 6 degree-of-freedom (DoF) change in position and attitude of the spacecraft. In the case of image-based motion estimation, the motion computed is the change in position and attitude between image captures. Image-based motion estimation has a long history in the machine vision literature, and the algorithm we use falls in the category of two-frame feature-based motion estimation. Our algorithm works by tracking multiple image features between a pair of images from which the spacecraft motion between the images is computed. Figure 2 describes pictorially the processes that occur during motion estimation. Below we give a brief overview of our motion estimation algorithm because it is a basis for current work. For more details, please see our previous work [5].

The first step in two-frame motion estimation is the extraction of features from the first image. Features are pixel locations and the surrounding image intensity neighborhood (call this a feature window) that can be tracked well across multiple images that may under go arbitrary, but small, changes in illumination or viewing direction. A qualitative definition of a good feature is a feature window that has strong texture variations in all directions. Since the motion between images is small, the change in position of features from image to image will be small. After feature detection, the features selected in the first frame are located in the second frame; this procedure is called feature tracking.

Feature detection and tracking have been studied extensively and multiple proven algorithms exist. Since processing speed is an important design constraint for our application, we selected the efficient feature detection algorithm of Benedetti and Perona [1] and the standard Shi-Tomasi feature tracker [10].

The motion between two camera views is described by a rigid transformation *(R, **T**)* where *R* encodes the rotation between views and ***T*** encodes the translation between views. Once features are tracked between images, the motion of the camera can be estimated by solving for the motion parameters that, when applied to the features in the first image, bring them close to the corresponding features in the second image

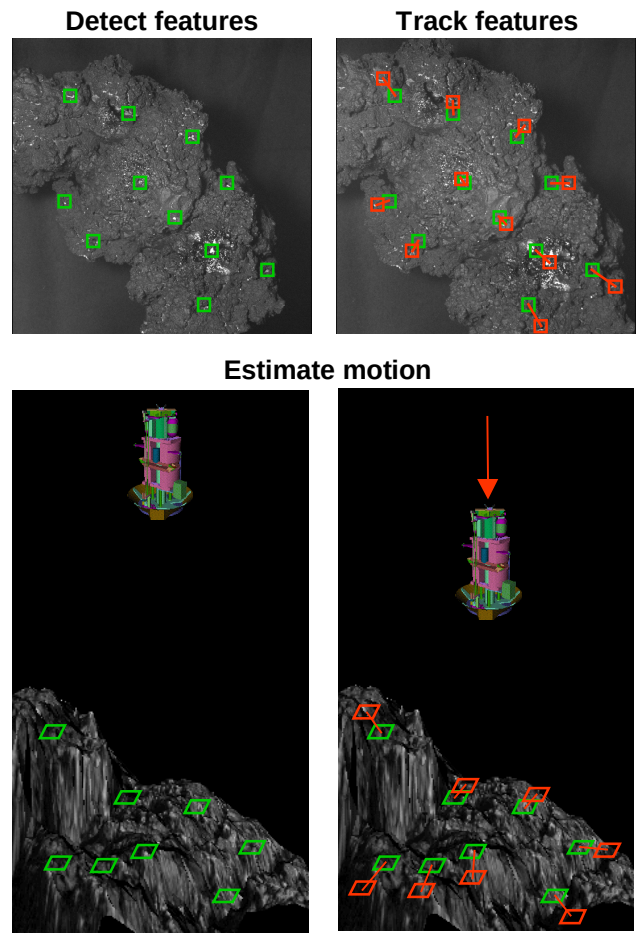

Figure 2 Feature-based motion estimation.

A fundamental shortcoming of all image-based motion estimation algorithms is the inability to solve for the magnitude of translational motion. Intuitively the reason for this is that the algorithms cannot differentiate between a very large object that is far from the camera or a small object that is close to the camera. Consequently, the output of motion estimation is a 5 DoF motion composed of a unit vector describing the direction of heading and the rotation matrix *R* between views. As described in [5] laser altimetry can be combined with 5 DoF motion estimation to compute the complete 6 DoF motion of the spacecraft. Other alternatives are to use on-board inertial measurement sensors or radiometric tracking from earth.

In our algorithm, motion estimation is a two-stage process. First an initial estimate of the motion is computed using a linear algorithm [7]. This algorithm is applied multiple times using different sets of features to eliminate feature track outliers and determine a robust Least Median Squares estimate of motion. The result of this algorithm is then used as input to a more accurate nonlinear algorithm that solves for the motion parameters directly. Since an good initial estimate is needed to initialize any nonlinear feature-based motion estimation algorithm, this two-stage approach is common [11]. Output from the nonlinear algorithm is the estimate of the five motion parameters and their covariance. This is then combined with laser altimetry or inertial measurements to create a 6 DoF motion estimate. For the lab imagery, the magnitude of translation is acquired from reading the translation stage used to move the camera during image stream acquisition.

It should be noted that all of our algorithms use a CAHVOR camera model [3] to intrinsically calibrate the camera. This camera model removes imaging nonlinearities due to radial lens distortion and optical center offset so that they do not effect motion estimation or structure recovery.

Once motion is computed, the 3-D position of the tracked features can also be computed using triangulation and the motion between the images. Below we show how the 3-D position of many tracked features can be used to reconstruct the surface topography, which can subsequently be used for hazard detection.

On tests conducted using real imagery we achieved a motion estimation rate of 4 Hz (on a 176 MHz R10000 processor) and motion estimation errors of less than 1% of the distance traveled. Using Monte Carlo simulation, we have shown that using only image-based motion estimation, it is possible to obtain a 3.6 m landing error ellipse when starting from a known position at an altitude of 1000 m.

Motion Stereo

Motion stereo is used to generate dense topographic maps of the imaged scene for use in absolute position estimation and hazard detection. First image-based motion estimation as described above is applied to determine the spacecraft motion between two images. To obtain a large baseline (distance between image acquisitions), features may be tracked through multiple images. When the desired baseline is reached, the motion is computed between the images based on the feature tracks using the algorithm described above.

If pixel matches between images are guaranteed to lie along the same scan line then a 1-D search, instead of a 2-D search, can be used to find pixel matches and a more efficient stereo matching algorithm can be realized. However, if the sensor rotates or translates out of the image plane between images then pixel matches will not lie along the scan lines of the images. In this case, image rectification is used to enforce the scan line alignment.

Image rectification creates two new images where pixel matches are guaranteed to lie along scan lines by creating two perfect virtual cameras that “view” the acquired images. The virtual cameras are created so that they have the same image plane, and they are oriented and positioned in the image plane so that corresponding rows in the virtual cameras define the same plane in space. The image plane and baseline between virtual cameras are chosen so that the rectified images are as close to the original images as possible. Projecting the original image pixels into the corresponding virtual camera then creates each rectified image.

To reduce the effect of intensity biases in pixel matching due to changes in illumination, a Laplacian filter (modeled as a difference of gaussians) is applied to the rectified images. This band pass filter eliminates the DC intensity and the noisy high frequency components of the images.

Next, matches at every pixel with sufficient texture are found along scan lines of the rectified and filtered images using a sum of absolute differences stereo matching algorithm [12]. This algorithm is highly optimized and has been applied to many mobile robot navigation problems. Once the pixel matches are computed, the depth to each pixel is computing using the motion between images and triangulation.

During small body exploration, the spacecraft will make motions that are beyond the scope of our binocular stereo vision algorithms (e.g., large rotations, translations out of the image plane). However with a few modifications, we were able to make our binocular stereo visions algorithms, work for motion stereo as well. The first change was to modify the rectification algorithm so that the rectified images were cropped to contain only the part of the image that contains data projected from the original images. With this change it was then necessary to change the stereo matching code so that it could search in both directions along scan lines for the best matching pixel. Although seemingly minor modifications, these changes allowed us to increase the amount of surface reconstructed in each image without increasing the running times of the algorithms.

Figure 3 shows the stages in the motion stereo algorithm for a wide baseline data set taken of a comet analog. The motion between the images is 10° about the vertical axis and a translation of 25 cm. Using this motion, the images are rectified, cropped and filtered as described above. Finally, stereo matching is performed and a dense depth map is constructed. The entire process of rectification through stereo matching took 4.5 seconds for 512x512 images on a 300Mhz Sparc Ultra 10 processor. The depth map displayed is color coded according to the visible spectrum. Red data is the closest to the sensor while magenta data is the farthest from the sensor. As the figure shows, very detailed surface topography can be generated using motion stereo and the

compute time is small enough to make motion stereo feasible during small body exploration.

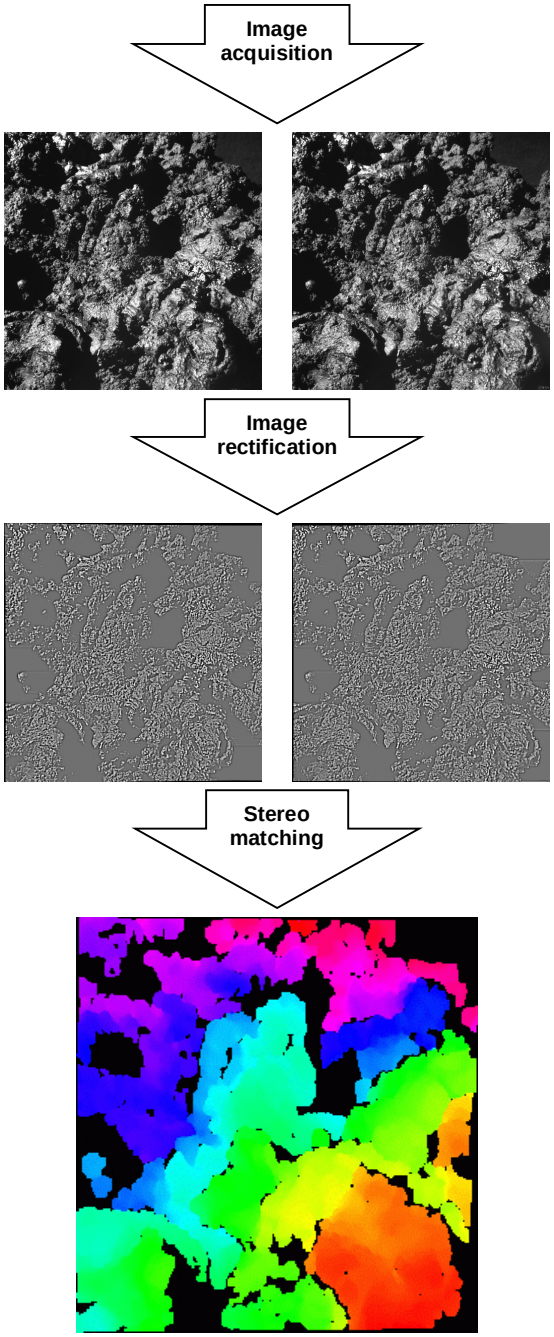

Figure 3 Motion stereo procedure.

Position Estimation

Image-based absolute position estimation is the process by which the 6 DoF absolute position of the spacecraft relative to a proximal small body coordinate system is determined from camera imagery. Absolute position estimation has two applications. Should the spacecraft lose all position and attitude information due to an anomaly, it can be used to realign the spacecraft with the small body coordinate system. The other more common application is to use absolute position estimation to null out dead reckoning errors from motion estimation during small body exploration.

We have taken the landmark matching approach to image-based absolute position estimation. First, a 3-D database of the small body landmarks is constructed from orbital imagery either automatically or on the ground with human intervention. During exploration, imagery is acquired and landmarks are extracted automatically. These landmarks are then matched to the 3-D landmark database and the position of the spacecraft in the small body frame is computed.

For landmarks to be effective, they must exhibit some invariance to illumination conditions and viewing direction. The general appearance of asteroids is well known from the multiple images taken during asteroid fly-bys; asteroids are marked with craters. In another paper [6], we show that craters can be used as effective landmarks for asteroid absolute position estimation.

However, in the case of comets, we have no high-resolution imagery of the comet surface, so establishing a strong geometric model for landmarks (e.g., craters are circular) is not possible. For comets, a more general landmark model must be used. There are three factors that indicate that surface shape can be used directly to describe landmarks for comet absolute position estimation. First, comet surfaces are expected to be rough on all scales, which makes surface shape very descriptive for matching. Second, although the appearance of a surface changes with illumination, the shape of the surface is invariant to illumination conditions. Third, there exist shape representations that provide local descriptions of surface shape that are invariant to surface position and attitude. These three factors make shape an attractive basis for landmarks. Since motion stereo can be used to extract the shape of a surface from monocular image streams, it is feasible to use a local shape representation to generate landmarks for comet absolute position estimation. Figure 4 through Figure 6 show that surface shape works well for absolute position estimation.

To represent landmarks, we use the spin-image shape representation [4], a local surface shape descriptor developed for surface alignment and object recognition. In this representation, surface shape is described by a dense collection of oriented points, 3-D points with surface normal. Using a single point basis constructed from an oriented point (point with surface normal), the position of other points on the surface can be described by two parameters. The accumulation of these parameters for many points on the surface results in an image at each oriented point. These images, localized 2-D patterns describing the local shape of the surface, are invariant to rigid transformations. Through correlation of images, point correspondences between two surfaces can be established in much the same way that templates are matched in 2-D computer vision. When two surfaces have many point

correspondences, they match. From this surface match, the absolute position of the spacecraft can be determined.

Spin-images from corresponding points on two different views of the same object will be similar, so spin-images can be used to establish point correspondences between surfaces. In this application, we match a 3-D model of the comet surface to a 3-D surface patch extracted from imagery using motion stereo. Briefly, the procedure is as follows; for more details consult [4]. In the initial offline stage, the 3-D model of the comet surface is constructed from orbital imagery, stored as a polygonal mesh, and the spin-images at each vertex in the mesh are generated. To initiate comet absolute position estimation, a sequence of images is acquired and the viewed surface patch is reconstructed using motion stereo. Next, a subset of vertices is selected from the surface patch and the spin-images for these vertices are generated. The spin-images from the surface patch are then compared to the spin-images from the model. When two spin-images are highly correlated, a point correspondence between the surface patch and the 3-D comet model is established. Point correspondences are then grouped and outliers are eliminated using geometric consistency. Groups of geometrically consistent correspondences are then used to calculate a rigid transformation that aligns the surface patch with the comet model. Finally, the alignment of the surfaces is verified and refined using a modified iterative closest point algorithm [13]. Since the surface patch is represented in the coordinate frame of the spacecraft, the transformation that aligns the surfaces also describes the absolute position of the spacecraft in the coordinate system of the comet.

To verify our comet absolute position estimation procedure we conducted three tests using laboratory imagery of a comet analog.. The analog was developed by a comet scientist at JPL to replicate our best knowledge of the appearance of a comet; it is rough on all scales and matte black with an albedo of roughly 4%.

As shown in Figure 4 the first test verifies the ability to match a small surface patch to a complete comet model. First a sequence of orbital images was taken in the lab by placing the comet analog on a turntable and rotating it in front of the camera. Features were then tracked through these sequences and the 3-D positions of the features were computed. These features were connected into a triangular mesh using a 2-D Delaunay triangulation of feature positions projected onto a cylinder defined by the orbit of the camera. This mesh constitutes the 3-D comet model. Next, a different set of images from the same orbit was taken. Features were tracked in these images and the corresponding structure was computed. These features were linked together into a triangular mesh using a Delaunay triangulation of the feature image positions. This mesh constituted the surface patch. The surface patch was then synthetically perturbed to an assumed position far from its correct location. Spin-images were then used to align the surface patch to the 3-D model from this assumed position.

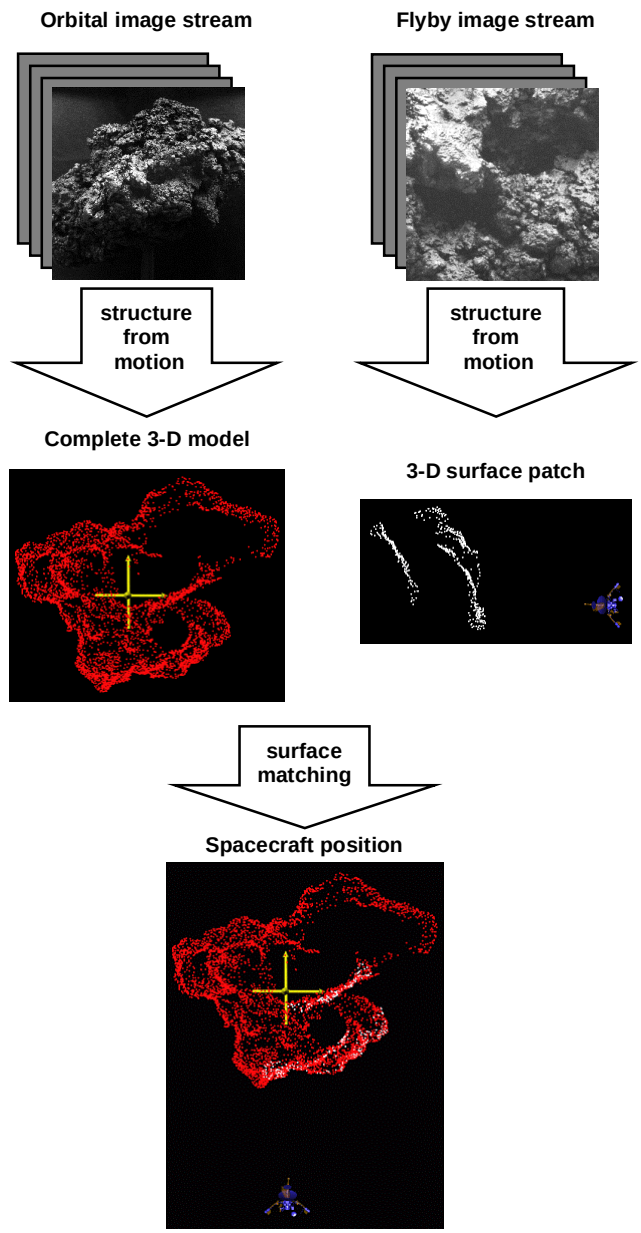

Figure 4 Absolute position estimation procedure.

Table 1 Position information for Figure 4

	Initial	Estimate	Truth
t_x (cm)	10.00	0.23	0
t_y (cm)	10.00	0.05	0
t_z (cm)	10.00	0.00	0
δ_t (cm)	17.32	0.24	NA
r_x (°)	50.00	0.41	0
r_y (°)	30.00	1.65	0
r_z (°)	50.00	0.59	0
δ_r (°)	76.81	1.80	NA

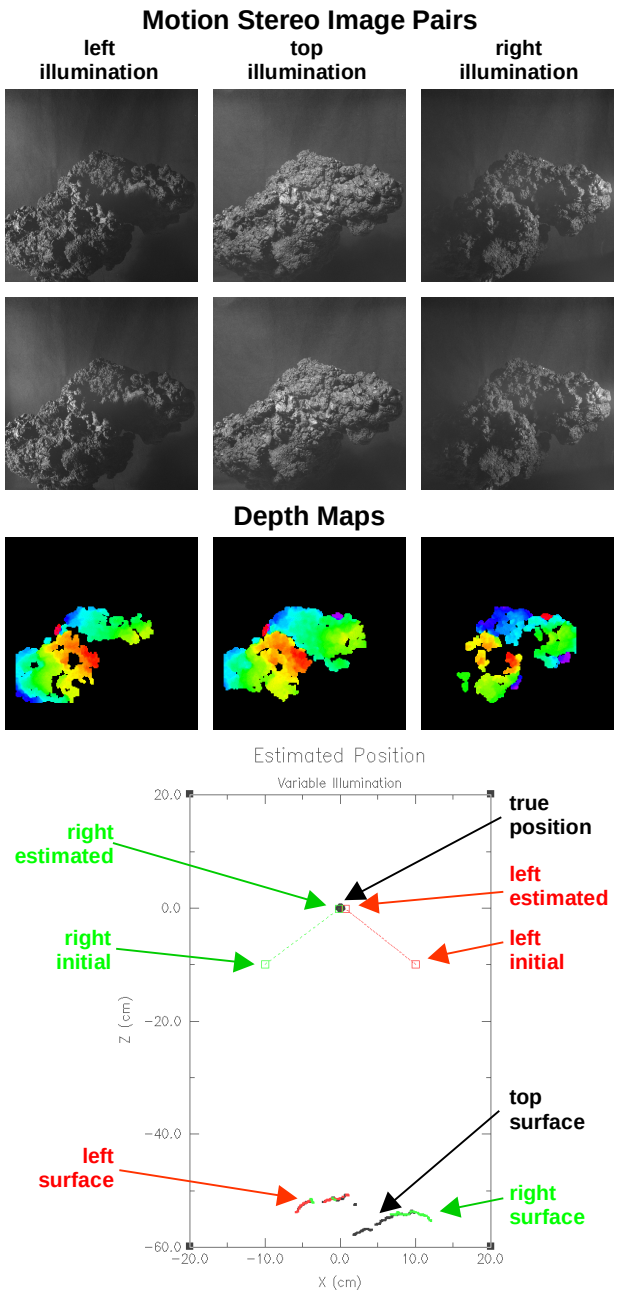

Figure 5 Illumination invariant position estimation.

Table 2 Position information for Figure 5

	Right			Left		
	Init.	Est.	Truth	Init.	Est.	Truth
t_x (cm)	-10.00	-0.17	0	10.00	0.68	0
t_y (cm)	10.00	0.11	0	10.00	0.34	0
t_z (cm)	10.00	0.07	0	10.00	0.10	0
δ_t (cm)	17.32	0.22	NA	17.32	0.77	NA
r_x (°)	50.00	0.13	0	50	0.36	0
r_y (°)	30.00	0.20	0	30	-0.67	0
r_z (°)	50.00	0.48	0	50	0.20	0
δ_r (°)	76.81	0.54	NA	76.81	0.79	NA

As shown in Figure 4 the surface alignment is quite good. Table 1 gives a comparison of positions and attitude, represented by Euler fixed angles, in order to assess the performance of surface matching. The assumed position is far from the true position with a Root Mean Square (RMS) error in translation of δ_t=17.32 cm and a RMS error in attitude of δ_r= 76.81°. However, the absolute position after surface alignment is dramatically improved with δ_t=0.24 cm and a RMS error in attitude of δ_r= 1.80°. This result verifies our approach to absolute position estimation.

The next test verified the performance of shape-based position estimation under variable illumination conditions. As shown in Figure 5 , the comet analog was placed in front of the camera and imaged under three different illumination directions. For each illumination direction a sequence of images was taken along a trajectory that caused the camera to rotate about a point on the comet surface. This trajectory included rotation and translation components and was the same for each illumination direction. The top of Figure 5 shows the first and last images in the sequences for a top, right and left illumination directions. Although the images are taken from the same position, the appearance of the images varies drastically due to the changes in illumination condition and subsequent shadowing.

The depth maps in the middle of Figure 5 were computed using motion stereo. The coverage of the depth maps varies based on illumination because you cannot reconstruct depth in the shadow regions due to the absence of texture. However, the depth values in the few regions of overlap between depth maps show similar depth variations, which demonstrates the surface shape can be reconstructed in a manner that is independent of illumination conditions.

We matched the right illuminated surface to the top illuminated surface and also matched the left illuminated surface to the top illuminated surface. At the bottom of Figure 5 we show the aligned surfaces and the initial and estimated positions of the sensor. Initially the sensor is assumed to be far from its true position. After alignment, the sensor position is much closer to its true position. This is shown qualitatively in Figure 5 (using a 2-D slice through XZ plane for clarity) and quantitatively in Table 2.

The final test verified the performance of shape-based position estimation under variable viewing trajectories. As shown in Figure 6 , the comet analog was placed close to the camera and three different sequences were taken. In the first sequence the camera was translated 0.25 cm and rotated 0.5° about the vertical axis between each image. The second sequence was constructed by tilting the camera 5° and then acquiring a sequence with 0.25 cm translation and 0.5° vertical rotation between images. The final sequence was acquired in a similar fashion except with a tilt of 10°. The top of Figure 6 shows the first and last images for each sequence. The depth maps in the middle of Figure 6 were computed using motion stereo.

The tilting of the camera is obvious from the orientation of the depth maps. We matched separately the 5 tilted and the 10 tilted depth maps to the 0 tilted depth map. The 2-D plot on the bottom left of Figure 6 shows the surfaces and sensor positions before alignment and the 2-D plot on the bottom right shows the sensor positions and surfaces after alignment. Although no ground truth position is available, the precise surface alignment indicates that the position estimation is correct. The timing for all 3 absolute position estimation experiments including surface patch spin-image generation, spin-image matching and pose estimation was less than 15 s on a 176 MHz R10000 processor. The positional accuracies as well as the rapid execution times demonstrate that shape, in the form of spin-images, can be used to generate effective view and illumination invariant landmarks for comet absolute position estimation.

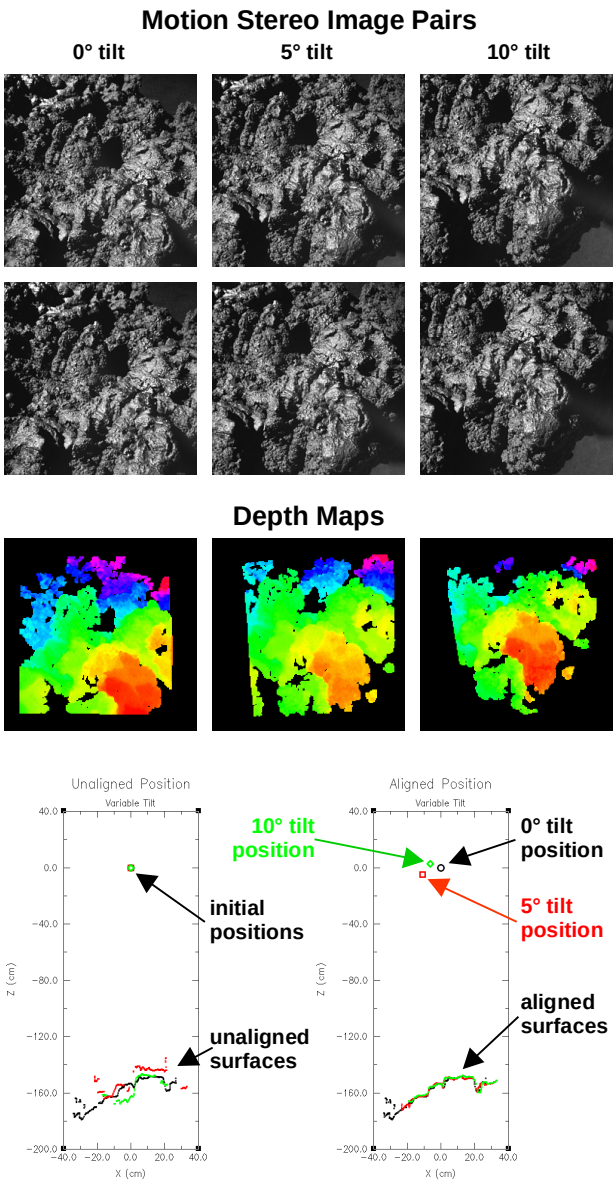

Figure 6 View-invariant absolute position estimation.

Hazard Detection

Hazard detection algorithms locate landing hazards in imagery while hazard avoidance algorithms guide the spacecraft away from hazards once they are detected. To date we have only investigated the hazard detection problem. Hazard detection depends on the mission scenario and the design of the spacecraft. As a baseline, we are using the ST4/Champollion spacecraft and mission scenario when designing our algorithms. In their scenario, regions of large slope and rough surfaces constitute hazards to the spacecraft. It is expected that by eliminating regions using these low-level hazards, that the high-level hazards such as crevasses and boulders will be detected and avoided. We can compute surface slope and roughness using the dense surfaces reconstructed using motion stereo. By applying constraints on surface roughness and slope, we can find the areas in the surface being imaged that are free of hazards.

We define surface slope at each pixel in a depth map by fitting a plane to the surface data in a local area around the pixel. The size of this area is set to the expected size of the lander footprint. The plane is fit using a standard least squares solution for plane fitting which has a closed form solution [2]. Given this local plane the slope of the surface with respect to a specified approach direction can be computed. By placing a maximum allowable slope constraint on this data, surface regions that are too oblique with respect to the landing direction are be eliminated.

Surface roughness is defined as the maximum absolute deviation from the best-fit plane of the surface data in a local area around a pixel. Once again this area is set to the expected footprint of the lander. This definition of roughness is appropriate because it will detect both rocks and crevasses. By applying a maximum surface roughness constraint, flat regions that exhibit too much surface variation for safe landing are eliminated.

We can measure the surface slope and roughness using the dense surface reconstructed using motion stereo. Constraints on maximum roughness and slope will be used to detect parts of the scene to be avoided, and given these constraints safe landing sites can be determined.

Two tests were conducted to assess the performance of hazard detection. In the first test, hazards were detected from same set of images used in Figure 6 . In these images, the camera was tilted between each sequence. Figure 7 shows an image for reference and the reconstructed terrain maps, for each sequence. Below the terrain map are images that show the roughness hazards (1 cm), the slope hazards (60°) and the safe zones for landing. The safe zones for all three sequences roughly correspond to the same places on the surface of the comet analog showing that our hazard detection algorithms have some view invariance.

In the second test, the hazards detected in a descent sequence were assessed. For descent imaging, the image epipoles lie in the image, making image rectification and

subsequent stereo matching impossible. However, a coarse terrain map can be reconstructed from the 3-D position of tracked features by projecting the depth of the features back into the image. Using this terrain map, hazards can be assessed in descent imagery.

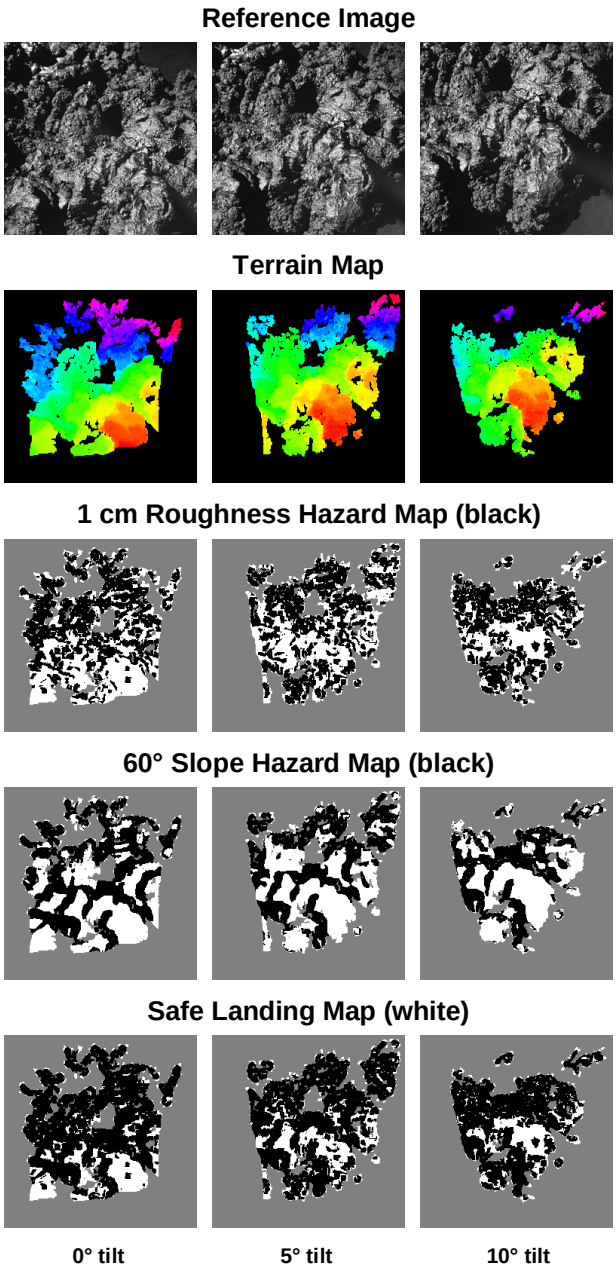

Figure 7 Motion stereo based hazard detection.

The top of Figure 8 shows the first and last image of a descent sequence. Below that are shown the tracks of the ¼ of the features used to generate the terrain map. Below the feature tracks are shown the reconstructed terrain map, the map of safe landing zones, the map of roughness hazards (1 cm) and the map of slope hazards (75°). The hole in the terrain map and subsequent maps indicates the area in the image where depth cannot be computed reliably because feature disparity is too small. Qualitative comparison of hazards to the terrain map shows that hazards have been correctly detected. This test shows that surface structure can be computed from descent sequences and that subsequently hazards can be detected using our algorithms.

4. Conclusions

The algorithms we have discussed are advancing the state of the art in vision-based navigation for small body exploration in many directions. First, an integrated set of algorithms for 6 degree of freedom motion and position estimation for comets and asteroids during descent and landing has never been developed. Second, these algorithms perform completely autonomously. These algorithms also provide techniques for dense surface reconstruction from monocular image streams that enable hazard avoidance and 3-D mapping for in-situ science. Finally, these algorithms are based on a single camera; this has positive implications in terms of power, cost and mass for any spacecraft utilizing these techniques.

The primary user of this technology will be small body missions. A future small body mission that can benefit greatly from our algorithms is Comet Nucleus Sample Return (CNSR). A requirement of CNSR is precision guidance and landing with hazard avoidance to three pre-determined sites on a comet nucleus. Other relevant future small body missions are Large Asteroid Sample Return, Asteroid Tomography and the Multi-Asteroid Trojans Flyby missions. This technology is also applicable to hazard avoidance during landing for the Europa Lander mission and 3-D surface mapping by aerobots during the Titan Organics Explorer mission.

In the near term, these algorithms can also be applied to the problem of safe and precise landing on Mars. Image-based motion estimation can provide estimates of rates on spacecraft attitude and position that can be used to control the trajectory of a lander during precision landing. Furthermore, during the terminal phase of descent, our algorithms for surface reconstruction from motion stereo and hazard detection can be applied to detect surface hazards and ultimately enable safe landing.

Although we have developed many algorithms for the small body navigation problem, there is still significant work to be done. Currently we are working on algorithms that tightly couple motion and position estimation so that high accuracy motion and position estimates can be obtained for longer periods. Better motion estimates will improve the accuracy of surfaces generated from motion stereo and will ultimately lead to more accurate hazard detection. Another area of work will be to modify existing algorithms to generate denser depth estimates from descent imagery. This will provide us with the data needed to match surfaces and consequently estimate absolute position from descent imagery. Another area of improvement will be to extend our motion and motion stereo algorithms to operate on multiple

image frames. In the end, we hope to demonstrate our algorithms in a real-time hardware-in-the-loop test on an unmanned aerial vehicle.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We would like to thank Jean-Yves Bouguet, Peitro Perona, Mark Maimone and Yalin Xiong for discussions and assistance leading to the results presented in this paper. We would like to thank Jackie Green for providing the comet analog.

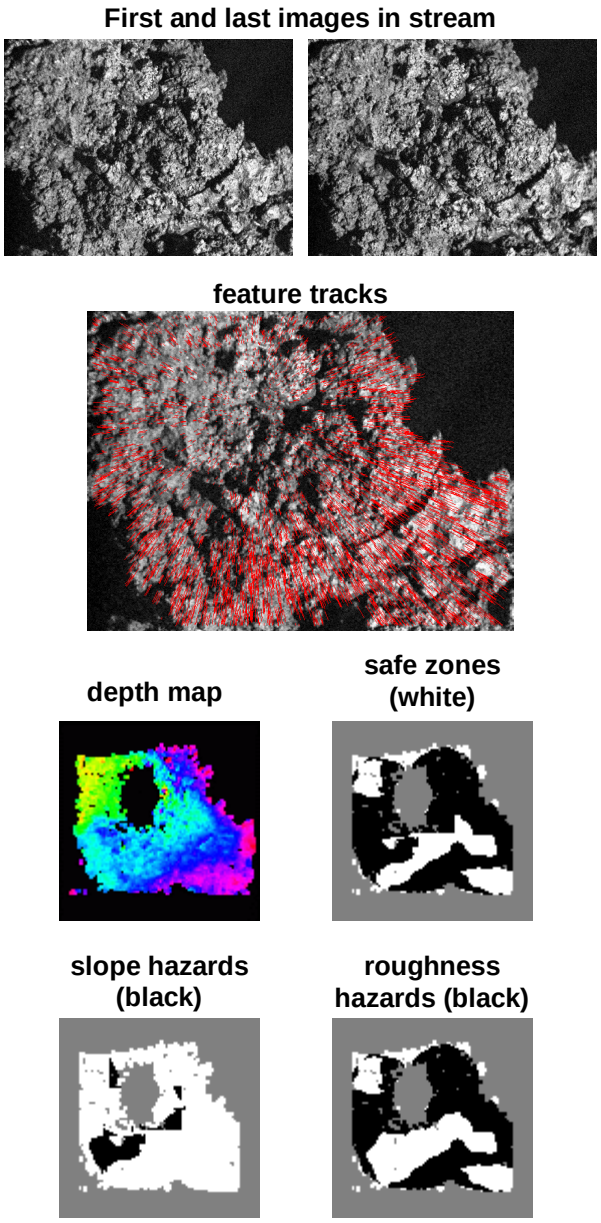

Figure 8 Hazard detection from descent imagery.

REFERENCES

[1] A. Benedetti and P. Perona, "Real-Time 2-D Feature Detection on a Reconfigurable Computer," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 586-593, 1998.

[2] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*. New York: Wiley-Interscience, 1973.

[3] D. Gennery, "Least-Squares Camera Calibration Including Lens Distortion and Automatic Editing of Calibration Points," To appear in *Calibration and Orientation of Cameras in Computer Vision*, A Grun and T. Huang, eds., Springer-Verlag, 1999.

[4] A. E. Johnson and M. Hebert, "Surface Matching for Object Recognition in Complex Three-Dimensional Scenes," *Image and Vision Computing* **16**, 635-651, 1998.

[5] A. E. Johnson and L. H. Matthies, "Precise Image-Based Motion Estimation for Autonomous Small Body Exploration," *Proc. 5th Int'l Symp. On Artificial Intelligence, Robotics and Automation in Space*, 627-634, June 1999.

[6] B. Leroy, G. Medioni, A. Johnson and L. Matthies, "Crater Detection for Autonomous Landing On Asteroids," *Proc. IEEE Workshop on Perception for Mobile Agents*, June 1999.

[7] H. Longuet-Higgins, "A Computer Algorithm for Reconstructing a Scene from Two Projections," *Nature* **293**, 133-135, September 1981.

[8] J. Miller et al., "Navigation Analysis For Eros Rendezvous And Orbital Phases," *Journal Astronautical Sciences* **43**(4), 453-476, 1995.

[9] J. Reidel, S. Bhaskaran, S. Synott, W. Bollman and G. Null, "An Autonomous Optical Navigation and Control System for Interplanetary Exploration Missions," *Proc. 2nd IAA Int'l Conf. on Low-Cost Planetary Missions*, IAA-L-506, 1996.

[10] J. Shi and C. Tomasi, "Good Features to Track," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 593- 600, 1994.

[11] J. Weng, N. Ahuja and T. Huang, "Optimal Motion and Structure Estimation," *IEEE Pattern Analysis and Machine Intelligence* **15**(9), 864-884, 1993.

[12] Y. Xiong and L. Matthies, "Error Analysis of a Real-Time Stereo System," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1087-93, 1997.

[13] Z. Zhang. "Iterative Point Matching for Registration of Free-Form Curves and Surfaces." *IJCV*, 13 (2), pp. 116-152, 1994

Andrew E. Johnson graduated with Highest Distinction from the University of Kansas in 1991 with a BS in Engineering Physics and a BS in Mathematics. In 1997, he received his Ph.D. from the Robotics Institute at Carnegie Mellon University where he studied three-dimensional object representation and recognition. Currently, he is a Senior Member of Technical Staff at the Jet Propulsion Laboratory where he researching image-based techniques for autonomous navigation and mapping around comets and asteroids. In addition, he is developing range image based precision guidance algorithms for the Comet Nucleus Sample Return pre-project. His general research interests are 3-D object recognition, surface registration, multi-view integration, environment modeling, 3-D structure and motion recovery from image streams, and autonomous navigation.

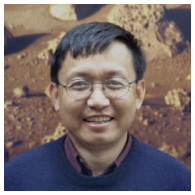

Yang Cheng received the B.S. and M.E. degrees from the Surveying and Mapping Institute, Zhengzhou, China and a Ph.D. degree in Remote Sensing at the University of South Carolina in 1993. During 1994 - 1999, he was a research staff member at Oak Ridge National Laboratory. Currently, he is a senior staff member at Jet Propulsion Laboratory. His major research interests include image processing, computer vision, remote sensing, image compression, cartography, map projection, parallel computation, and Geographic Information Systems.

Larry H. Matthies obtained a Ph.D. degree in Computer Science from Carnegie Mellon University in 1989, then moved to the Jet Propulsion Laboratory, where he is currently a Principal Member of Technical Staff and supervisor of the Machine Vision Group. His research has focused on terrain sensing and obstacle avoidance algorithms for autonomous navigation of robotic vehicles. At JPL, he pioneered the development of real-time algorithms for stereo vision-based obstacle detection and he contributed to the development of the structured light sensor used by the Sojourner Mars rover. His group currently has research projects on computer vision for robotic vehicles sponsored by NASA, DARPA, and the U.S. Army; these projects include work on navigation of Mars rovers, asteroid and comet landers, and Earth-based robotic vehicles for urban and cross-country missions. He is a member of the editorial board of the Autonomous Robots journal and an adjunct member of the Computer Science Department at the University of Southern California.