# Terrain Tracking for Lander Guidance Using Binary Phase-Only Spatial Filters

*Max B. Reid* and *Butler P. Hine*
Computational Systems Research Branch, NASA Ames Research Center, Mail Stop 269-3, Moffett Field, CA 94035

ABSTRACT: We demonstrate the use of binary phase-only spatial filters (BPOFs) for tracking target sites through a sequence of gray-scale terrain images. The filters are demonstrated in a closed loop guidance system for a laboratory lander mockup. Images of a 3D terrain board taken by the lander's video camera are preprocessed to produce binary intensity contour maps of a simulated planetary surface. A BPOF is made from a section of the current preprocessed image centered on the exact desired landing site. After the lander has descended to a lower altitude, the BPOF is correlated with a new image. The position of the correlation peak is used in making the next filter and to guide the lander so as to recenter the landing site in the camera's view. We present testbed results of the accuracy with which a site may be tracked from orbit to landing, and the maximum scale, translation, and rotation which can be tolerated between subsequent images. Application of the results to NASA's proposed Mars Rover Sample Return (MRSR) and Mars Environmental Survey (MESUR) missions is discussed.

## INTRODUCTION

THE SUCCESS OF NASA's future autonomous planetary landing missions depends on the accurate and safe landing on a site selected from an initial image taken from orbit. Current techniques (Vaughan *et al.*, 1991; Gaskell, 1988) for autonomous planetary lander guidance involve terrain following optical navigation. Images taken from a lander camera during descent are preprocessed to isolate distinctive surface features, or landmarks. Each image is compared to images in a precomputed library of landmark images expected to be seen at different points along the lander's trajectory. In this way, the lander's actual position is compared to its expected position at a given time.

Spatial filtering has been applied to processing terrain imagery (Pan, 1989), but has not been applied in practice for planetary lander guidance. Efforts to track locations in terrain images through correlation have been limited to small kernal correlations wherein specific features isolated through preprocessing are tracked from image to image. Direct correlation of large terrain images with full amplitude and phase matched filters has not been employed because of the poor quality of results achieved. Figure 1a shows a 256 level gray-scale image of desert terrain in a 128- by 128-pixel field. There are few distinct features in the image, and when a matched filter is made from the central region of the image, Figure 1b, the correlation of the original image with this filter, Figure 1c, does not yield an accurate identification of the target site's location.

## BPOFs FOR TERRAIN SITE CORRELATION

In other target recognition applications, binary phase-only filters (BPOFs) have been found to produce correlations with certain advantages over matched filters (Horner and Leger, 1985; Psaltis *et al.*, 1984). BPOFs produce correlation peaks which are much more localized than those produced by matched filters. However, the tradeoff for this sharper correlation peak is increased sensitivity to geometrical distortions and to noise. A BPOF designed to recognize an object scene, $s(x,y)$, is made from the object's Fourier transform, $S(u,v)$, through the equation

$$H(u,v) = \begin{cases} 1 & Re[S(u,v)e^{i\phi}] > 0, \\ -1 & Re[S(u,v)e^{i\phi}] < 0. \end{cases} \quad (1)$$

The filter is binarized about the angle $\phi$ in the complex plane, as shown in Figure 2. Most commonly, binarization is made about the imaginary axis with $\phi = 90°$, which is equivalent to saying that binarization is based on the sign of the real part of $S(u,v)$ (Downie and Reid, 1990). For isolated objects, such as the binary silhouette of a Space Shuttle Orbiter shown in Figure 3a, filtering the object with a BPOF made from its Fourier transform yields an extremely sharp correlation peak. The BPOF made from the Shuttle using Equation 1 with $\phi = 90°$ is shown in Figure 3b, and the resulting correlation of the image and filter is depicted in Figure 3c. In contrast, the correlation of the image

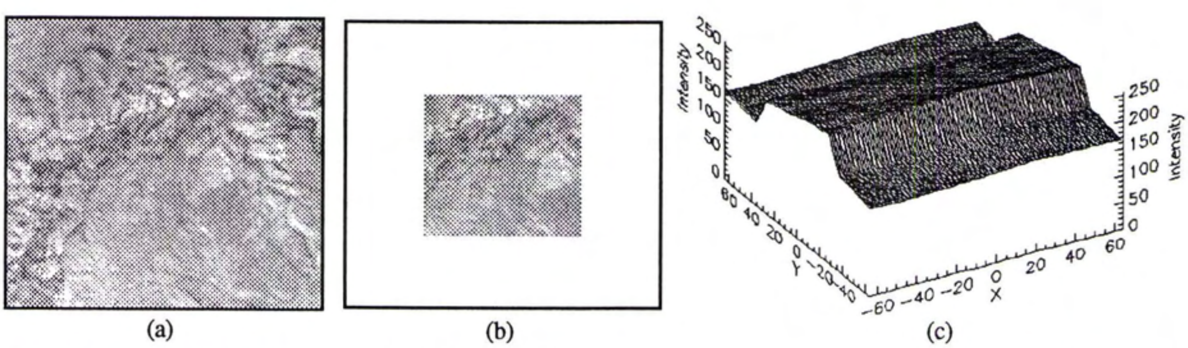


(a) (b) (c)

FIG. 1. Gray-scale terrain image correlation. (a) 256-level image of desert terrain in 128- by 128-pixel field. (b) Central 64- by 64-pixel region of image from which matched filter is made. (c) Intensity of correlation of filter with original image.

PHOTOGRAMMETRIC ENGINEERING & REMOTE SENSING,
Vol. 58, No. 12, December 1992, pp. 1699–1706.

0099-1112/92/5810–1699$03.00/0

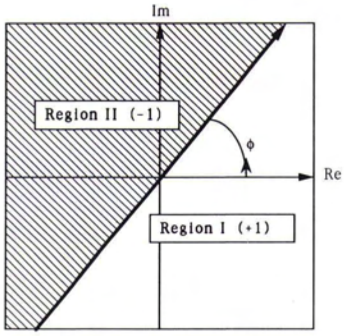


FIG. 2. Complex plane with real and imaginary axes and binarization angle, $\phi$.

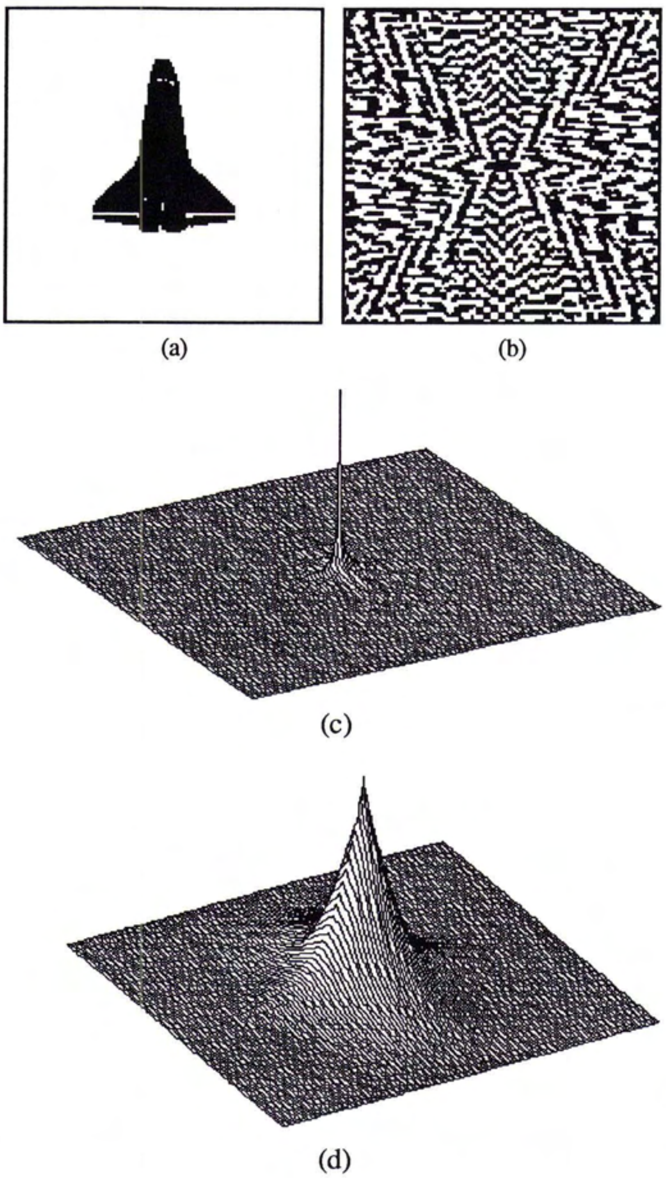


FIG. 3. Binary phase-only correlation. (a) Space Shuttle Orbiter. b) BPOF made from Oribter with $\phi = 90°$. (c) Correlation of Orbiter with BPOF. (d) Correlation of Orbiter with its fully complex matched filter.

with its (unbinarized) matched filter is shown in Figure 3d. The correlation peak for the BPOF is much more localized, allowing more accurate determination of the target's location.

In this paper, we attempt to determine the performance capabilities of BPOFs for the purpose of tracking sites in terrain images. The term "site" here refers to any sub-region of an original gray-scale terrain image, for instance, a central region as in Figure 1. A BPOF made from the site in Figure 1b is shown in Figure 4a and the resulting correlation of this filter with the original image in Figure 1a is shown in Figure 4b. The result is far from the single clean correlation peak seen in Figure 3c for the Shuttle Orbiter. The reason is that the gray-scale image is relatively featureless, and can be described as a constant background intensity level with a slowly varying modulation superposed. A filter made from any sub-region of such an image is to a large extent just a filter for the window aperturing the sub-region, and it will produce a relatively high correlation peak with every sub-region of the original image which is the same size as the aperturing window. The smaller the target site window, the less the target area can be discriminated.

We have seen in Figure 4 that, even with a target site size as large as 64 by 64 pixels, we achieve relatively poor correlation performance due to the featureless nature of the imagery. If we desire to use a target site substantially smaller than the entire image to allow significant translation of the observing vehicle between images, preprocessing of input images must be performed to enhance features. We need a fast preprocessing algorithm which does not involve time consuming image segmentation. In the algorithm used here, an original image, $I$, is converted to a binary filled-contour map, $C$; i.e.,

$$C(x,y) = \begin{cases} 0 & I(x,y) \text{ modulo } (512/N) < 256/N, \\ 1 & \text{otherwise,} \end{cases} \tag{2}$$

where $I(x,y)$ is the image intensity prescaled so that $I(x,y)_{max} = 255$. For example, Figure 5b is made from the gray-scale image in Figure 5a using Equation 2 with $N = 4$. In this example, all values of $I(x,y)$ between 0 and 63 are mapped to 0, intensities between 64 and 127 are mapped to 1, 128 through 191 to 0, and 192 through 255 to 1. Figures 5c and 5d show two other possible mappings with $N = 2$ and $N = 8$. The choice of $N$ for best terrain tracking performance is discussed later.

To see the effect of this input image binarization algorithm, consider Figure 6. The mapped image in Figure 6a was made from Figure 1a using $N = 4$ in Equation 2, while the BPOF in Figure 6c was made from the 64- by 64-pixel centered target site in Figure 6b. The correlation of this filter with the full contour image (Figure 6a) is shown in Figure 6d. The result now is a single distinct correlation peak. Obtaining this distinct peak required both binarizing the input image and using a binary filter. For comparison, in Figure 6e we show the correlation of the binarized image of Figure 6a with a fully complex matched filter made from the target in Figure 6b. While the correlation does have a clearly recognizable peak, as opposed to the matched filter correlation with the original gray-scale image shown earlier in Figure 1c, the correlation plane is much more cluttered and the peak much less localized than when a binary phase-only filter is used.

## TERRAIN TRACKING EXPERIMENTAL PROCEDURE

The utility of BPOF correlation for terrain tracking is tested by demonstrating closed loop guidance of a laboratory lander mockup. The lander consists of a Microbot laboratory robot controlled by a Sun Microsystems SPARC workstation. Mounted on the robot is the 0.1-kg camera head of a Panasonic WV-CD1BW camera. The camera looks down on a three-dimensional terrain board created from Landsat images of the Grand Canyon region in Arizona. A

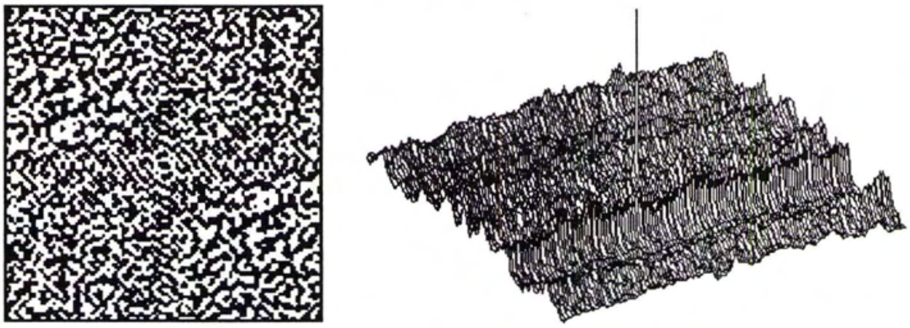

FIG. 4. Terrain site correlation with BPOFs. (a) BPOF made from gray-scale terrain site (Figure 1b). (b) Intensity of correlation of gray-scale terrain image (Figure 1a) with BPOF made from terrain site.

photograph of the test station is shown in Figure 7. With the iris of the 7.5-mm focal length, $f$1.6 lens closed to a pinhole aperture, the depth of field of the camera is sufficient to focus images of the terrain board at heights between 40 and 200 mm.

To test the lander guidance system, the lander is positioned at an altitude of 200 mm above the 3D board. The landing sequence then follows the algorithm depicted in Figure 8. A landing site is selected from an image taken at this initial "on-orbit" position by a person using a mouse. The lander is then automatically controlled to keep the landing site centered while the lander descends to an altitude of 40 mm. Figure 9a shows an initial gray-scale image at 200 mm with the site selection crosshair. Figure 9b shows the camera's view after a descent sequence down to 40 mm. Figure 9 shows that an image of the landing site taken at lowest altitude contains significantly more detail than was apparent at the beginning of the landing sequence. The terrain board is built on a scale of 100,000 : 1, so that 200 mm corresponds to an altitude of 20 km and 40 mm corresponds to 4 km. However, the exact altitude represented by the camera's height above the terrain board is not critical. We are interested in determining the ability of the tracking algorithm to remain locked on a target site while the lander's altitude decreases by a significant factor, in this case a factor of five.

Images are taken and preprocessed with an Androx ICS-100 frame grabber to produce 128- by 128-pixel binary intensity contour maps of the simulated planetary surface. A BPOF is made from a 64- by 64-pixel section of the current preprocessed image centered on the exact desired landing site. After the lander has descended to a lower altitude, the BPOF is correlated with a new (preprocessed) image. The position of the correlation peak is used to make the next filter and to guide the lander so as to recenter the landing site in the camera's view.

The parameters which we desire to measure are the accuracy with which a site can be tracked during descent, and the maximum scale, translation, and rotation which can be tolerated between subsequent images. For the purpose of these experiments, the landing sequence is deemed to succeed when each correlation produces a peak at a location which is within two pixels of the actual target site location. Binary phase-only filters are much more sensitive to scale and rotation distortions than are fully complex matched filters; this is a tradeoff for the much more localized correlation peak obtained with BPOFs. The first case tested is tolerance to scale changes alone, with no rotation or translation during descent. At each step, the lander descends by a constant factor of its current altitude, producing constant scale changes between consecutive images. For instance, for a descent factor of 0.04, the lander's altitude is decreased by a total factor of 5 geometrically by steps equal to 4 percent of the current altitude, requiring a total of 41 steps, each step including a correlation and the making of a new filter.

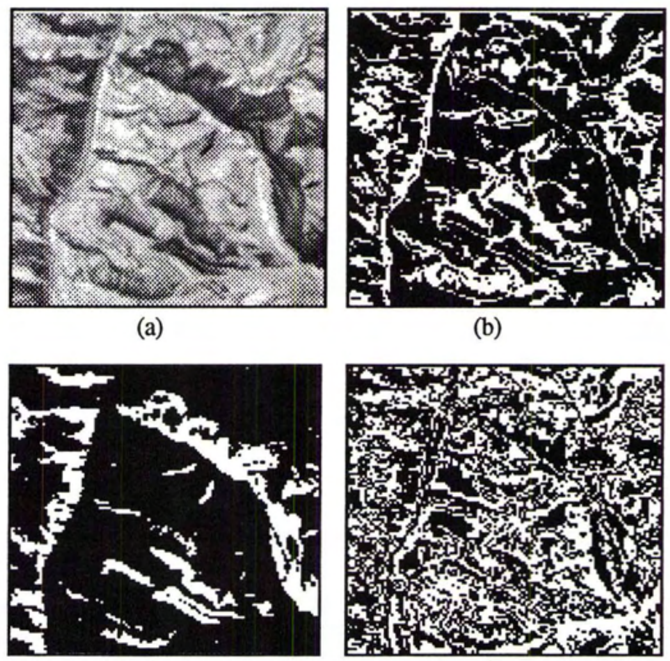

FIG. 5. Binary filled-contour image mapping. (a) Original gray-scale terrain image. (b) Binary mapping of original image with four contour levels. (c) Binary contour mapping with two levels. (d) Binary contour mapping with eight levels.

The next cases tested are tolerance to in-plane rotation and to translation, both at a fixed altitude. Spatial correlations should be invariant to translation. The intensity peak in the correlation plane moves as the image moves in the input plane, and thus the location of the correlation peak provides a measure of the location of the input image. However, the field distortion of the small fast camera lens limits the amount of translation which can be tolerated in the test system. After determining the tolerance of the BPOF correlation guidance system to each of scale, rotation, and translation of the input scene individually, the case of simultaneous distortions is considered. This simulates

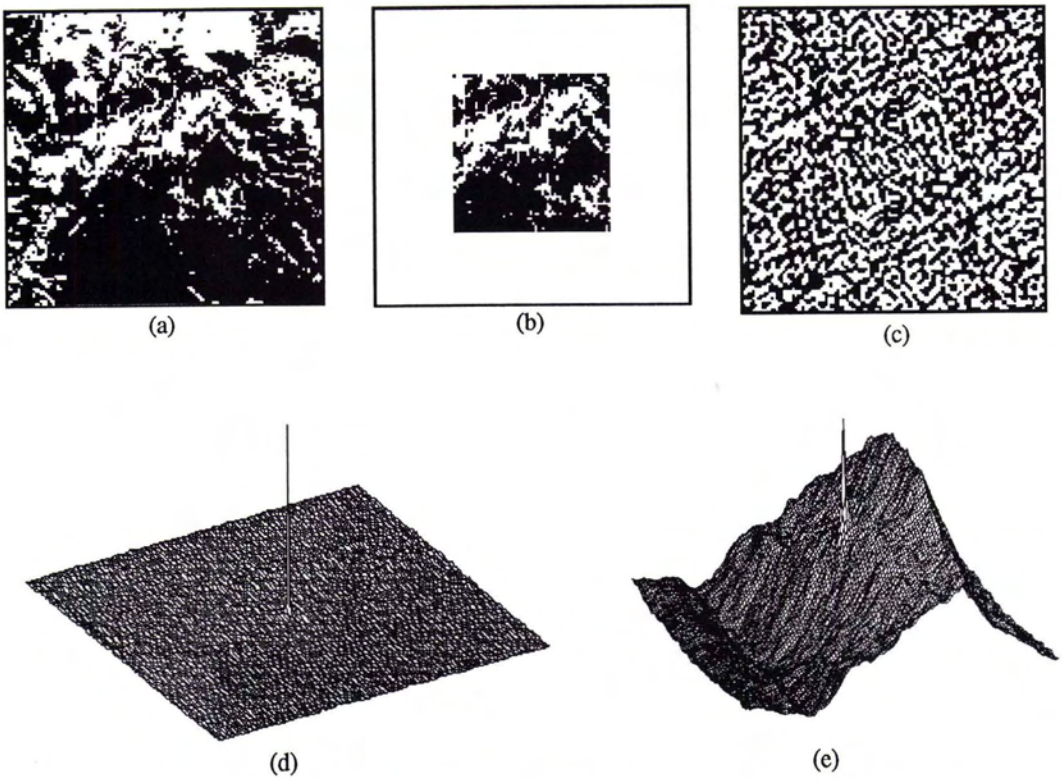


FIG. 6. Terrain site correlation with binary contour images and BPOFs. (a) Binary filled-contour mapping of original gray-scale terrain image (Figure 1a). (b) Target site consisting of 64- by 64-pixel region from the center of the mapped image. (c) BPOF made from binary target site. (d) Intensity of correlation of target site filter with full 128- by 128-contour image. (e) Intensity of correlation of a fully complex matched filter made from the target site with the full 128- by 128-contour image. The peak is much less localized than in (d).

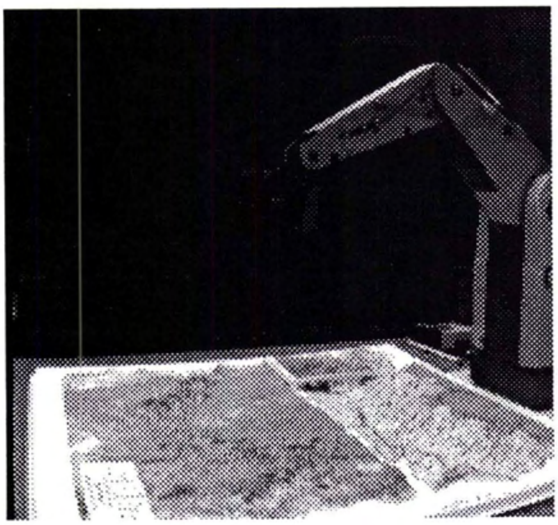

FIG. 7. Photograph of lander guidance test station. Microbot robot positions Panasonic WV-CD1BW camera over a three-dimensional terrain board created from Landsat images of the Grand Canyon region in Arizona.

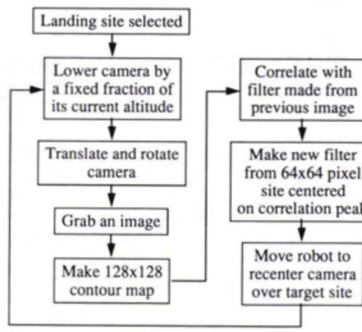


FIG. 8 . Algorithm for testing lander vision system based on binary phase-only filters.

the case where the landing vehicle is driven slightly off course, for instance by Martian winds, and is rotating during descent.

In the preceding cases, it is assumed that no preprocessing is performed on input images to compensate for scale and rotation. In some lander scenarios, instruments may provide data on vehicle altitude and angular orientation. To measure tolerance to scale and rotation in this case, we either pre-scale or pre-rotate each new image by the amount the lander has descended or rotated since the last filter was made. The greatest tolerable scale or rotation in this case is mostly a function of the accuracy of the vehicle's altitude and orientation measurements. This experiment, using an inexpensive laboratory robot without its own position sensors, provides an example of a BPOF guidance system's performance for a lander with relatively inaccurate altitude and orientation knowledge. The final case tested includes both rotation and translation of the lander during descent, where input images are pre-scaled and rotated before being correlated with the current filter.

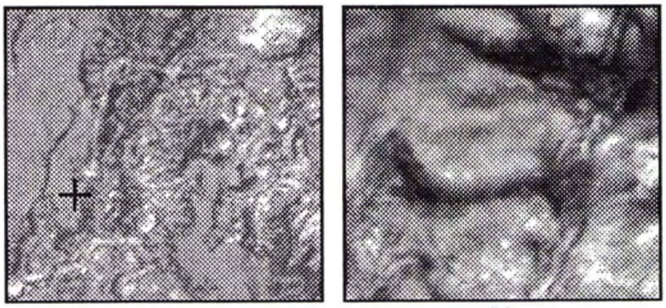

(a) (b)

FIG. 9. Typical test station images. (a) Gray-scale 128- by 128-pixel terrain image taken from lander mockup camera at 200-mm height above terrain board. (b) Terrain board image at 40 mm.

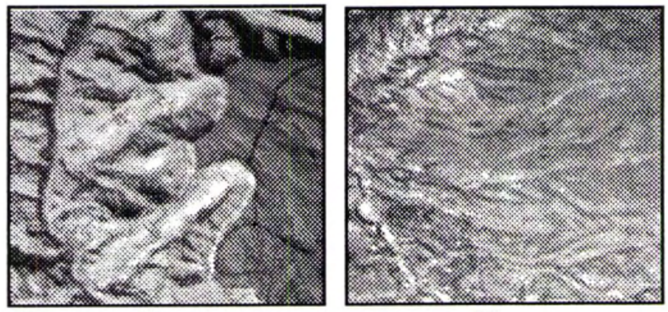

(a) (b)

FIG. 10. Test landing sites. (a) Cliff area with distinct terrain features. (b) Plateau region with few terrain features.

## EXPERIMENTAL LANDER RESULTS

### NUMBER OF CONTOUR LEVELS

During initial terrain tracking experiments to determine tolerance to geometric distortions, we also performed tests to determine the optimal number of contour levels to which to map gray-scale input scenes. Best performance was found with four levels, as shown in the example in Figure 5b and given by Equation 2. Mappings with fewer levels, such as in Figure 5c, tended to fail at low altitude when the landing site was in a particularly flat, featureless desert area. This is attributable to the fact that the 64- by 64-pixel target area would often become nearly completely uniform, either all ones or zeros. In contrast, more contour levels, such as in Figure 5d, led to image mappings with extremely high frequency features. This tended to make the system very sensitive to scale and rotation changes during descent.

### SIZE OF TARGET AREA

We discussed above the choice of landing site target areas which were chosen as 64- by 64-pixel regions centered on the landing site and cut out of the 128- by 128-pixel contour image. Smaller site sizes produced correlations with lower signal to noise, and this made the system less tolerant of geometric variance between images. The 64- by 64-pixel site size was chosen as it gives relatively high correlation peaks while still allowing 32 pixels of translation between input scenes before the target begins to be apertured.

### TOLERANCE TO SCALE CHANGES

The descent sequence was performed as described earlier, first with very small incremental descent factors, then with ever larger factors until tracking failed. Tests were conducted by descending on each of two different target areas, an area with fairly distinct features, the cliff area in Figure 10a, and an area with very few features, the plateau in Figure 10b. Tracking failed when any of the correlations made during this process failed to have a peak within 2 pixels of the tracking site.

The greatest incremental reduction in lander altitude which allowed repeatable successful tracking of both sites was 4 percent. Descending from 200 mm to 40 mm geometrically by a factor of 0.04 requires 41 steps in the landing sequence. As noted, the location of the target site obtained from the correlation peak was no more than two pixels off from the correct position in any of these 41 steps. A typical three-dimensional plot of a correlation peak with a 4 percent scale change between filter and target in shown in Figure 11a. The peak is clearly discernible, even though it is reduced in intensity by a factor of 6 from the peak intensity of the correlation of the filter with the unscaled image, as seen in Figure 11b which shows the

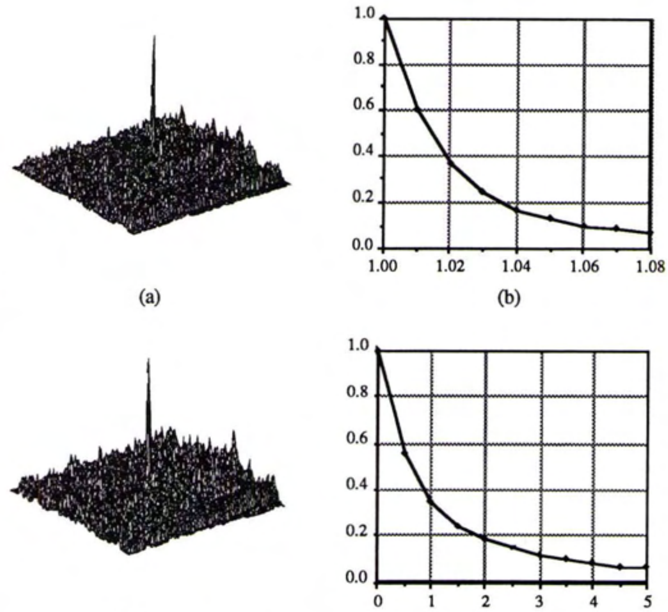


(a) (b) (c) (d)

FIG. 11. Degradation of BPOF correlation peaks due to scale and in-plane rotation distortions. The undistorted image is the contour map shown in Figure 6a. Original gray-scale image in Figure 1a is scaled or rotated and then binarized to form distorted images. The filter is made from the target in Figure 6b. (a) Correlation of image increased in size by 4 percent with undistorted filter. For comparison, correlation with undistorted image is seen in Figure 6d. (b) Plot of correlation peak intensity vs. scale distortion of image; 0 to 8 percent scale change. (c) Correlation of image rotated by 2.5° with undistorted filter. (d) Plot of correlation peak intensity vs. in-plane rotation distortion of image; 0 to 5° rotation.

sensitivity of terrain site correlation to scale changes between the input image and the image from which the target filter is made.

### TOLERANCE TO TRANSLATION

A guidance system based on correlation should be invariant to translation in the transverse plane up to the point where the target site no longer appears fully in the input scene. This would be 32 pixels of displacement in either transverse direction, which for our optics corresponds to a distance equal to 14 percent of the lander's height above the terrain board. The field distortion produced by the camera used in this experiment limited the tolerable translation to ≈10 pixels, or about 5 percent of the lander altitude. This limitation is an artifact of the current experiment, and the larger value of 32 pixels of tolerable translation will be used below in evaluating the applicability of BPOF correlation to planetary lander guidance.

### TOLERANCE TO ROTATION

To test the sensitivity of a BPOF correlation guidance system to rotations between successive input scenes, the lander was held at a constant altitude and rotated by 90° and back by increments of increasing size. Tracking again failed when any of the correlations made during this process failed to have a peak within 2 pixels of the tracking site. The largest in-plane rotation which could be tolerated was 2.5°. The system sucessfully and repeatably tracked targets for each of 72 steps while rotating at this rate.

Figure 11c shows that a typical correlation between a filter and an image rotated by 2.5° has a very distinct intensity peak, even though it is reduced in by a factor of 6 from the peak intensity of the correlation of the filter with the unrotated image, the same reduction as for a 4 percent scale change, as seen above. Figure 11d shows the sensitivity of terrain site correlation to in-plane rotation between the input image and the image from which the target filter is made.

### SIMULTANEOUS SCALE, TRANSLATION, AND ROTATION

Rather than test every possible combination of scale, translation, and rotation distortion, tests were started with the three distortions set at the maximum level found tolerable for each individually (scale: 4 percent of altitude, translation: 5 percent of altitude, rotation: 2.5°). All three distortions were then decreased proportionately until the system was able to repeatably track both the landing sites shown in Figure 10. The maximum tolerable simultaneous distortions were equal to ≈60 percent of the values tolerable for each distortion individually. Rotations of 1.5° and translations of 3 percent of altitude were tolerated while descending by steps equal to 2.5 percent of altitude.

### TOLERANCE TO SCALE (COMPENSATED)

If a landing system includes instruments for determining the craft's altitude, then images input to a correlation guidance system can be pre-scaled to compensate for altitude changes between images. To test this case, before input images were binarized to form filled-contour maps, they were scaled down in size by the same factor by which the altitude was decreased at each step. The result was successful tracking up to scale changes of 16 percent. Only 11 steps of this size are required to descend by a total factor of 5. Prescaling therefore increases the time available to complete each correlation cycle by approximately a factor of 4, as the prescaling operation itself is much faster (by approximately $\ln(128^2) = 14$) than a correlation / filter creation sequence.

The value found here for the increase in tolerable scale changes obtainable using prescaling is fairly conservative. The degree to which compensation can increase tolerance to scale changes depends on how accurately the altitude is known. Our inexpensive laboratory robot has no position sensors, so the accuracy with which it moves from a point A to a point B depends strongly on the accuracy of the system's estimate of where point A is. As many movements are made, the robot controller's estimate of the robot's altitude becomes progressively more inaccurate. Using prescaling, a landing system with active altitude sensing should be able to tolerate greater scale changes between images than the value of 16 percent found here.

### TOLERANCE TO ROTATION (COMPENSATED)

If a lander has on-board gyroscopes to measure its orientation, then input scenes may by pre-rotated to remove distortions between successive images. In this experiment, we were able to tolerate in-plane rotations up to 20° when input scenes were pre-rotated prior to correlation, an increase in rotation invariance of a factor of 8. Again, the limit on this technique is determined more by the accuracy of the orientation measurement than by the inherent capabilities of BPOF correlation, and the value of 20° is a conservative estimate for the amount of in-plane rotation between successive images which a landing system with active gyroscopes should be able to tolerate while maintaining closed-loop guidance.

### SIMULTANEOUS SCALE, TRANSLATION, AND ROTATION (COMPENSATED)

Because our limited translation invariance is an artifact of the compact lander mockup optics, translation was not compensated. However, it was included in this test in order to demonstrate tracking through a realistic sequence of descent images. Scale and rotation were compensated. The three distortions were first set at their individual tolerable limits (scale: 16 percent of altitude, translation: 5 percent of altitude, rotation: 20°) and decreased proportionately until repeatably successful landing on the two sites of Figure 10 was achieved. The maximum tolerable simultaneous distortions were again reduced to about 60 percent of the values tolerable for each distortion taken singly. Rotations of 12° and translations of 3 percent of altitude were tolerated while descending by steps equal to 9 percent of altitude.

### LANDING ACCURACY

A landing sequence was deemed to have failed if any correlation step missed the actual landing site by more than 2 pixels. During successful landings, most correlation steps either tracked the site exactly or missed by only one pixel. In addition, the tracking errors were not systematic, but random. The lander descended in a random walk about the target site, though errors at higher altitudes had the most effect on the final outcome. In approximately two-thirds of successful tests, the system's estimate of the target location after reaching 40 mm altitude was no more than 2 mm in error, equivalent in our system to approximately a 2-pixel error in the initial image at 200 mm, or 1 percent of initial altitude. The tracking accuracy may be considered to be a normal random variable with $\sigma \approx 0.01\ A_0$, where $A_0$ is the altitude at which the tracking algorithm is initiated.

## APPLICATION OF BPOF CORRELATION TO PROPOSED MARS MISSIONS

Two future NASA missions are being studied which will involve the autonomous landing of remote probes on the surface of Mars. The Mars Rover Sample Return (MRSR) mission plans to soft-land a rather sophisticated lander which will include an autonomous roving vehicle to gather soil and rock samples over a wide area. The samples will be returned to Earth via a relaunch stage of the lander. The second mission under consideration is the Mars Environmental Survey (MESUR) which envisages hard-landing 20 very simple landers at scattered locations on the Martian surface. These landers will carry very little instrumentation and will utilize as few mechanisms as feasible to accomplish their mission of seismic and meterological data collection.

In this section, we extrapolate the performance of our laboratory mockup BPOF correlation based vision system to the conditions of these two proposed Mars missions. While these conditions are different in many ways from those of our lander mockup, we hope to achieve a rough estimate for the potential capabilities of BPOF correlation for lander guidance, particularly with respect to the accuracy of guidance and the update rate at which a BPOF correlation vision system must operate for planetary lander guidance.

### MRSR VISION

The MRSR lander will have thrusters, allowing the capability of altering its trajectory during descent. In order to land as close as possible to a landing site preselected from images taken from Mars orbit, a vision system is required to provide updated information to the trajectory control system concerning the vehicle's position with respect to the landing site. There are two critical issues of concern for the vision system: (1) how accurately can a target site be tracked, and (2) can the vision system provide updated information quickly enough to both remain locked on target and to close the trajectory control loop.

The lander mockup discussed above showed that the accuracy with which a vision system based on 128- by 128-pixel BPOF correlation can guide a lander to a preselected site is equal to approximately 1 percent of the initial distance from the lander to the target. Figure 12 shows the position of a lander with respect to a Martian landing site at the beginning of the approach sequence (Smith and Carter, 1989). The lander is at an altitude of 35 km and is 250 km uprange of the site. The lander is therefore 252 km from the site, with an angle of view of $\alpha = 41°$. A rough estimate (which will be modified by the exact details of the lander camera's field of view) of the expected landing footprint is given by the ellipse shown in Figure 12 with $a = 2.5$ km and $b = 2.5/\sin 41° = 3.8$ km. This footprint is roughly equivalent in size to the best performance expected from alternative vision-based guidance systems (Vaughan *et al.*, 1991).

The accuracy of the BPOF-based guidance system in these experiments was limited mostly by relatively small (1 or 2 pixel) errors in correlation peak locations. A single error of this size at highest altitude is alone nearly enough to cause an error in tracking equal to ≈1 percent of the initial altitude. This error may be reducible by using higher resolution input scenes and filters with 256 by 256 or more pixels.

Increased image/filter size will increase the time required to process each new image. To evaluate how quickly each step of calculating a correlation and making a new filter must be performed, we make the following assumptions: (1) feedback from the control and vision systems keep the lander's camera pointed directly at the landing site, the field distortion of the camera is negligible and does not limit translation invariance - therefore, 32 pixels of transverse displacement, equivalent to ≈15 percent of the current altitude, between consecutive images is tolerable; (2) on-board instrumentation provides altitude data, so scale changes between images can be pre-compensated, allowing scale changes of at least 9 percent; (3) vehicle orientation is measured, so rotations can be compensated, allowing in-plane rotation of at least 12° between images; and (4) out-of-plane rotation of an essentially flat image caused by changes in angle-of-view are equivalent to scale changes in one dimension and are tolerable to the same 9 percent level.

Invariance to translations of 15 percent of current altitude and in-plane rotations of 12° is more than enough to allow the vision system to measure expected trajectory errors. We therefore concentrate on scale changes caused by closing range and changes in angle-of-view. Using data from Smith and Carter (1989), dividing velocity by range at the beginning of landing approach gives a relative scale change of only 0.6 percent/sec. Given tolerance to 9 percent changes, 15 seconds are available in which to perform a correlation and create a new filter. At the end of initial approach and immediately before parachute deployment, the scale change rate has increased to 3 percent/sec, still allowing 3 seconds for update. The greatest scale changes in the entire landing sequence occur after parachute deployment and immediately before retrorocket propulsion ignition. This occurs at an altitude of 1.4 km, where the scale change rate is 8.6 percent/sec. During this portion of the landing sequence, the vision system update rate must reach nearly 1 Hz.

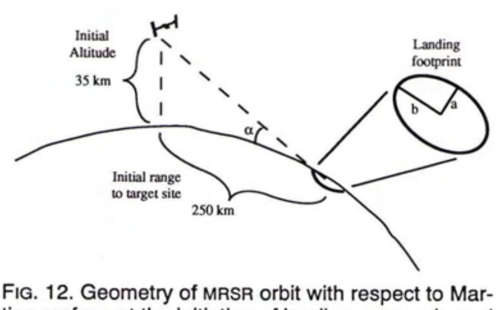


FIG. 12. Geometry of MRSR orbit with respect to Martian surface at the initiation of landing approach, and projected MRSR landing footprint with BPOF correlation based lander guidance.

Changes in angle-of-view during descent also do not require a high update rate. During approach, $\alpha$ changes by no more than ≈0.2°/sec, and the greatest scale change rate in one dimension caused by a change in angle-of-view is only ≈0.1 percent/sec. After parachute deployment and throughout the remainder of the landing sequence, the rate is even lower. We have neglected the possibility of relatively large out-of-plane rotations due to the lander swinging on its parachute. The degree of such rotation is difficult to determine at this time; however, as active guidance is not possible during the parachute deployment period of the landing sequence, rotation on the parachute will not significantly effect the size of the landing footprint.

The final consideration in determining the required response rate for the lander's vision system is the rate at which updated position information is required to close the trajectory control loop. Current studies show the required position update period does not need to be shorter then 2 or 3 seconds (H. Pien, personal communication, 1991). Therefore, the update rate required for the vision system to maintain tracking is more than sufficient to close the trajectory control loop, and the vision system never needs to be updated at greater than 1 Hz.

## MESUR Lander Vision

Mars Environmental Survey landers are projected to be much simpler than the MRSR lander (Hubbard *et al.*, 1991). A MESUR lander will not have active thrusters to control its trajectory during descent; therefore, landing footprint size cannot be affected by a terrain tracking vision system. However, measurement of the relative displacement between consecutive lander camera images taken during descent would be extremely important for this mission. MESUR landers will carry no instruments to measure altitude or either vertical or horizontal velocity. Figure 13 shows three different terminal landing conditions for a MESUR lander. Without altitude and velocity measurements, as in Figure 13a, the lander will not be able to use a retrorocket and will impact Mars at > 30 m/s. Using an airbag to cushion the impact, the deceleration will be approximately 40 Earth gravities. With knowledge of the velocity vector and altitude, the force the lander must be designed to withstand may be reduced by over an order of magnitude. In Figure 13b, a simple retrorocket can be used to greatly reduce the lander's vertical velocity if measurement of the lander's altitude is possible. Finally, if both altitude and velocity are measured, the lander's momentum, both vertically and horizontally, can be greatly reduced by aiming a "smart" retrorocket just before impact, as in Figure 13c.

A lander's altitude and velocity will be well known before it enters the Martian atmosphere. If terrain tracking is initiated before entry, the results of BPOF correlation can be used to maintain updated estimates of altitude and velocity throughout approach and parachute descent. A BPOF made from an image taken at altitude $A_1$ may be correlated with several scaled copies of an image taken at altitude $A_2$. The scale of the image which correlates best with the filter, combined with knowledge of $A_1$, yields an estimate for $A_2$ and vertical velocity. With an altitude

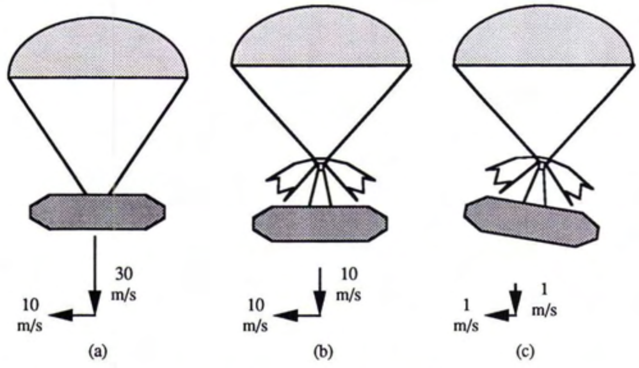


FIG. 13. Three landing conditions for MESUR landers. Vertical velocity on parachute is originally approximately 30 m/s. Horizontal velocity due to Martian winds has a one standard deviation value of 10 m/s. (a) No retrorocket: parachute all the way to the surface. (b) Simple retrorocket: fixed total impulse, altitude measurement only, preselected ignition height. (c) "Smart" retrorocket: variable impulse, vectored thrust, vertical and horizontal velocity estimation, real-time ignition command.

measurement in hand, the transverse displacement of a terrain site in consecutive images, measured by correlation in pixels, may be converted to a measurement of horizontal velocity.

Parachute deployment for a MESUR lander occurs at an altitude of 5 to 8 km. The vehicle's terminal vertical velocity after deployment is calculated from Mars atmosphere models to be ≤ 30 m/s, and descent will require 3 to 5 minutes. Initially, the scale of images taken from the lander's camera change at ≈0.5 percent/sec. There are, therefore, approximately 18 seconds available for altitude calculation. Even if each estimate requires making a BPOF and correlating it with several scaled copies of the next image, each correlation would not have to be performed faster than ≈3 seconds.

Approximately the first minute of parachute descent could be spent making several estimates of altitude and vertical velocity. Afterwards, with these parameters known, measurements of horizontal velocity could be updated in the time required to make one filter and perform one correlation. If this rate were 1 Hz, for instance, the final velocity estimate would be made at an altitude of 300 m. Below this, the scale change in images taken from the vehicle descending at 30 m/s would be too large to ensure accurate terrain tracking. At 300 m, the precision of transverse position measurement would be limited to ≈1.5 m by the size of the pixels in a 128 by 128 image field. To insure that horizontal velocity is known to 1 m/s or better accuracy, an update rate of ≈1.5 Hz would be necessary. Velocity measurements accurate to 1 m/s will allow a "smart" retrorocket to be employed to reduce the deceleration a lander must survive to only 1 to 2 Earth gravities, instead of the 40 gravities deceleration a lander without altitude and velocity measurement capabilities would have to survive (Figure 13). As for the MRSR lander, we have neglected the possibility of relatively large out-of-plane rotations due to the MESUR lander swinging on its parachute. This possibility will need to be thoroughly assessed in future studies.

## CONCLUSIONS

We have demonstrated on an experimental terrain board that BPOF correlation can be used as the basis of a lander guidance vision system. A landing site chosen from an initial terrain image can be tracked through a succession of images taken while a lander descends, rotates, and undergoes transverse displacement with respect to the landing site. If the lander's altitude and orientation are not known, scale changes between successive images are limited to 4 percent, and rotations to 2.5°. Even with relatively inaccurate measurement of altitude and orientation provided by on-board instrumentation, scale changes of at least 16 percent and rotation of more than 20° are tolerable. In both cases, these values are reduced by ≈40 percent if scale, rotation, and translation occur simultaneously.

BPOF correlation based vision could be used on both the Mars Rover Sample Return (MRSR) and the Mars Environmental Survey (MESUR) missions. We have extrapolated the results for the experimental mockup lander to obtain an estimate of the performance possible for these missions. For MRSR, a landing footprint of 8 km is obtainable with 128- by 128-pixel images and filters. This is comparable to the smallest footprint achievable through other tracking algorithms under consideration. The update rate required for each correlation/BPOF creation step is not required to be greater than 1 Hz. This is easily achievable with off-the-shelf digital signal processing (DSP) chips. If higher resolution images and filters are required to reduce the landing footprint further, custom digital electronic or analog optical systems (Huang *et al.*, 1991; Reid *et al.*, 1991) may be necessary.

For a MESUR lander, BPOF correlation with 128- by 128-pixel images and filters can provide a means for maintaining knowledge of the vehicle's altitude and velocity vector after atmospheric entry. With a correlation rate of 1 Hz, the final velocity measurement would be accurate to ≈1.5 m/s. This measurement would greatly reduce the force which the lander must survive on impact. The required update rate is again so low that the algorithm may be implemented on a standard DSP chip. A linear increase in the final velocity measurement can be achieved by increasing the update rate. Increasing the size of the image/filter field is not necessary. As the maximum update required will be at most a few hertz, custom hardware will not be required.

## ACKNOWLEDGMENT

This work was supported by the NASA Office of Aeronautics and Space Technology under RTOP 506-59-31.

## REFERENCES

Downie, J. D., and M. B. Reid, 1990. Mapping Considerations for Optimal Binary Correlation Filters, *Applied Optics*, 29:5235–5241.

Gaskell, R. W., 1988. Digital Identification of Cartographic Control Points, *Photogrammetric Engineering & Remote Sensing*, 54:723–727.

Horner, J. L., and J. R. Leger, 1985. Pattern Recognition with Binary Phase-Only Filters, *Applied Optics*, 24:609–611.

Huang, C. C., G. Gheen, and E. R. Washwell, 1991. Throughput Comparison of Optical and Digital Correlators for Image Recognition, *Proc. SPIE Conf. on Pattern Recognition V*, paper #1564-40, San Diego, 23-26 July.

Hubbard, S., *et al.*, 1991. Minutes of the Mars Environmental Survey Design Concept Review Meeting, NASA Ames Research Center, Mountain View, California, 1 May.

Pan, H. J., 1989. Spectral Analysis and Filtering Techniques in Digital Spatial Data Processing, *Photogrammetric Engineering & Remote Sensing*, 55:1203–1207.

Psaltis, D., E. G. Paek, and S. S. Venkatesh, 1984. Optical Image Correlation with a Binary Spatial Light Modulator, *Optical Engineering*, 23:698–704.

Reid, M. B., *et. al.*, 1991. Optical Computing at NASA Ames Research Center, *Proc. AIAA Computing in Aerospace 8*, Baltimore, 22-24 October, pp. 503–515.

Smith, R. S., and P. H. Carter, 1989. *MRSR Approach Through Landing Reference Trajectory*, Lockheed Engineering and Space Company Correspondence No. 6-89, Houston, Texas.

Vaughan, R. M., R. W. Gaskell, P. Halamek, A. R. Klumpp, and S. P. Synnott, 1991. Autonomous Precision Landing Using Terrain-Following Navigation, AAS/AIAA Spaceflight Mechanics Meeting, Paper AAS 91-137, Houston, Texas, 11-13 February.

(Received 6 December 1991; accepted 14 April 1992; revised 5 May 1992)