

Stereo Vision and Shadow Analysis for Landing Hazard Detection

Larry Matthies, Andres Huertas, Yang Cheng, and Andrew Johnson

Abstract— Unmanned planetary landers to date have landed “blind”, without the benefit of onboard landing hazard detection and avoidance systems. This constrains landing sites to very benign terrain and limits the scientific goals of missions. We review sensor options for landing hazard detection, then identify an approach based on stereo vision and shadow analysis that appears to address the broadest set of missions with the lowest cost. We describe algorithms for slope estimation and rock detection with this approach, develop models of their performance, and validate those models experimentally. Instantiating our model of rock detection reliability for Mars predicts that this approach would reduce the probability of failed landing by at least a factor of 4 compared to blind landing. Conversely, for the safety level desired for the 2009 Mars lander, this approach would increase the fraction of the planet that is accessible for landing from about 1/3 to nearly 100%.

I. INTRODUCTION

Landing site selection procedures for planetary exploration use all available remote sensing data to characterize the safety of potential landing sites before landing is attempted. With cameras now in orbit around Mars and planned to orbit Earth’s Moon, it is possible to map all landing hazards larger than a few meters across. However, descent navigation accuracy may not be sufficient to avoid small hazards seen from orbit. Moreover, slopes on the scale of a lander (e.g. < 6 m across) and rocks that could be fatal to a lander (eg. > 50 cm tall) may not be detectable from orbit. Many sites of scientific interest on Mars, in the lunar highlands, and on other moons and asteroids have rock distributions high enough to create a landing failure probability of several percent for blind landers. In contrast, the Mars Science Laboratory (MSL) lander/rover in development for a 2009 launch will accept a landing failure probability due to rock impalement of only 0.25%. For a blind landing, this rules out well over half the surface of the planet. Increasing the accessible surface area and reducing the probability of a failed landing requires onboard landing hazard detection (HD) and avoidance capabilities.

Sensor options for HD have been studied for many years, including lidar, radar, and passive imaging [1], [2], [3], [4]. Lidar and radar are attractive because they are direct ranging sensors applicable at relatively high altitudes. However, many factors make passive imaging attractive, including a shorter development cycle and potential for smaller size, lower power consumption, and lower cost [2]. Landers typi-

cally carry descent cameras for scientific imaging that could also be used for HD. A navigation camera may also be needed at high altitude for landmark recognition for precision navigation, which may also be used for HD.

There are many potential approaches to HD with passive imaging, including use of color, texture, shading, shadows, structure from motion (SFM), stereo, and visible vs. thermal spectral bands. The characteristics of the terrain, lighting, practical descent trajectories, and spacecraft payload constraints determine which approaches are feasible for any given mission. Selected methods also must have a statistical model of hazard detection performance, which has been validated experimentally, that can be used to show that the probability of landing failure is within acceptable limits.

Section II examines planetary landing scenarios to identify a set of sensor/algorithm alternatives with the broadest applicability and lowest anticipated cost, as well as to determine nominal sensor performance requirements. Our conclusion is that stereo vision and shadow analysis appear to cover the widest set of missions with the least cost. Section III summarizes algorithms we have developed to date for slope estimation and rock detection with these sensing modalities. Section IV describes performance models and experimental evaluation of these algorithms. Section V incorporates these results into an overall model of safe landing probability with these sensors. This model predicts that this approach would reduce the probability of a failed landing on Mars by more than a factor of 4 compared to a blind landing, assuming the ability to maneuver to avoid detected hazards. Conversely, for the 0.25% limit on rock impalement probability required by MSL, the model shows that this approach would nearly triple the fraction of the planet that would allow a safe landing. Much of our modeling is also applicable to hazard detection with lidar.

II. LANDING SCENARIO AND SENSOR OPTIONS

One of the most challenging places to do landing hazard detection is Mars, because its high gravity requires very fast HD and because its atmosphere reduces image contrast and constrains when sensing can be done. Therefore, we use Mars as a design driver, since methods that work there will apply to most target bodies.

The descent sequence designed for the upcoming MSL mission provides a well-defined reference scenario. The spacecraft enters the atmosphere at about 100 km above ground level (AGL). The heat shield is dropped and the parachute is deployed at about 10 km AGL. The parachute is released and retrorockets begin firing around 1.5 km AGL. A lateral divert maneuver begins at about 1.2 km AGL and

Manuscript received September 14, 2008. This work was supported by NASA under the New Millennium Program for the ST9 mission.

Authors are with the Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, CA 91109 USA (phone: 818-354-3722; fax: 818-354-9973; e-mail: firstname.lastname@jpl.nasa.gov).

ends at about 100 m AGL to get clear of the parachute. The horizontal movement during this maneuver covers about 25% of the starting altitude. Doing precise terrain relative navigation (TRN) by map matching before this divert would allow the divert to be targeted to avoid large hazards known from orbital reconnaissance, such as craters up to ~ 100 m in diameter. Detecting small scale hazards before or during this divert is impractical for several reasons:

- It would be expensive, because it would require very high sensor angular resolution over a wide field of regard.
- It would require very accurate navigation to guarantee avoiding all small scale hazards detected at more than 1 km away.
- During the divert, the high spacecraft attitude rates would make it difficult to obtain low smear, high SNR images aimed at the right place on the ground.

At the end of this divert, descent is vertical and relatively slow, so HD is then possible to enable a small divert of 1-2 lander diameters to avoid small-scale hazards, such as rocks. Thus, performing HD at or below ~ 100 m AGL appears to be most practical for MSL-like missions. Such a capability would also cover most other missions and most other target bodies [5].

Minimizing the size, weight, and power of the hazard detection sensors is desirable to minimize their cost. Lidar sensors under study for landing applications have mass, power, and volume estimates roughly in the 15-30 kg, 100-200 W, and 20,000 cm³ range for the sensor alone, not including data processing [3]. In contrast, visible spectrum descent cameras have been built on the scale of 0.6 kg, 10 W, and 490 cm³ [6]. Field programmable gate arrays (FPGA) can implement the computing required for hazard detection with a few kg and W. Therefore, passive imaging has potential to provide a smaller, lower power solution than lidar and radar, particularly if the same camera can also serve as a science camera or provide imagery for precision navigation before the start of descent.

Color, texture, and shape from shading with descent imagery are not promising for HD for a variety of reasons, including results from prior missions that show negligible color variation on asteroid Eros [7] and the impracticality of getting slope and rock size information with sufficient accuracy from texture and shading. Contrast in thermal imagery can discriminate rocks from soil over part of the diurnal cycle [8]. However, the vast majority and the highest resolution orbital mapping imagery is visible spectrum, so using thermal imagery for HD would increase cost by requiring separate visible and thermal cameras for TRN and HD.

Shadows can be used to recognize hazardous rocks from altitudes of 1 km or more [2], but this does not enable slope estimation. SFM can enable slope and rock detection if the descent trajectory can give adequate parallax and if the camera can be aimed at the landing site from two or more locations during the descent. This may be practical for missions

to comets, asteroids, and small moons, but it is costly and difficult for large bodies like Mars. Binocular stereo baselines of ~ 1 m or more appear to be feasible for most landers and can enable slope and rock detection at altitudes up to about 100 m. Given that this fits the challenging reference mission scenario above, stereo vision is our primary approach. Shadow analysis can enhance rock detection for small incremental runtime cost and can significantly increase rock detection altitude for missions where that is needed, so we include shadow-based rock detection in our approach. Based on our current knowledge of hazard densities around the solar system, this approach is applicable to most or all lander missions. As we discuss below, the speed, reliability, and hardware maturity of this approach makes it a candidate for missions in about four years.

There is still interest in lidar for HD, particularly for landing in shadowed regions of the lunar poles. However, lidar hardware is further from maturity for lander applications than descent cameras. Nevertheless, the HD algorithms and performance modeling we apply to range data from stereo are applicable to lidar as well.

III. HAZARD DETECTION ALGORITHMS

Planetary landing sites seen to date are fairly smooth surfaces with scattered rocks emplaced as impact crater ejecta [9]. Landing site safety requirements specify maximum slope angles, to avoid tip-over hazards, and maximum rock abundance, to avoid impaling the lander on a rock. Given this and constraints on the size and performance of space-qualifiable computers, we have developed simple algorithms for stereo-based slope estimation, stereo-based rock detection, and shadow-based rock detection. These algorithms are amenable to real-time implementation in space-qualifiable FPGAs and our experimental results show that they should provide satisfactory levels of detection reliability. More sophisticated algorithms are constrained by projections of the computing power likely to be available in future missions; this trade-off is a subject of ongoing research. This Section briefly summarizes the current state of these algorithms; Section IV evaluates their performance.

A. Stereo-based Slope Estimation

We employ a stereo vision algorithm based on real-time, area-based image cross correlation with the sum-of-absolute-differences (SAD) operator with either one correlation window per pixel (SAD1) [10] or extended to use five overlapping correlation windows per pixel (SAD5) [11]. We have implemented both algorithms in FPGA to generate subpixel disparity maps for different image sizes for different applications [12]; for SAD1 on 1024x768 pixel imagery, throughput is 15 frames/second (fps) using one Xilinx Virtex-II Pro FPGA. Fig. 1 illustrates stereo results from a data set used for evaluation in Section IV.

The slope estimation algorithm uses range images produced by stereo matching to compute a slope estimate by robust plane fitting. We first perform a least median squares

plane fit that includes all points in the image, including the rocks (outliers) on the surface. Next, we discard points far from the plane and apply a least squares fit to the remaining points to obtain the slope estimate. The algorithm has been tested with data that emulates “altitudes” up to 100 m by placing cameras at different distances from the wall shown in Fig. 1. Fig. 2 shows an example of plane fitting applied to the wall range data at 30 m and at 70 m. This shows that the plane fits are approximately correct even with rocks present as outliers. Further performance modeling and evaluation is discussed in Section IV-A.

We plan to implement the plane fit and the rest of the hazard detection algorithms described below in a PowerPC 405 processor embedded in the Virtex-II Pro FPGA, with floating point-intensive calculations potentially done by a 100 MHz PowerPC 750-based main spacecraft computer. We expect that the end-to-end process of stereo vision followed by rock and slope detection from range data and shadows will achieve a throughput of 1 fps or better in this processor configuration, which is fast enough for our reference mission scenario.

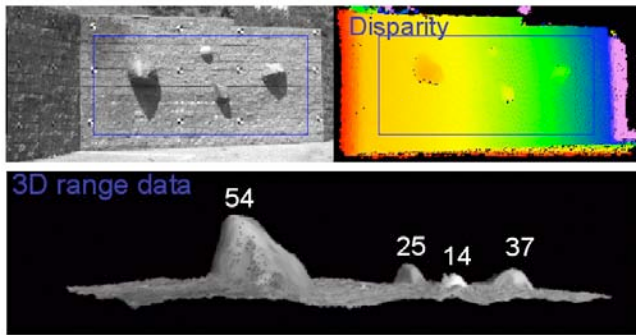


Fig. 1. Sample SAD5 stereo vision-based range imaging results. Upper left: $\sim 6\text{ m} \times 3\text{ m}$ brick wall with four synthetic rocks viewed from 40 m distance. Upper right: false color range image; red is closest and magenta is furthest. The overlaid rectangle shows the area used to evaluate plane fitting for slope estimation. Bottom: 3-D rendering from below. Numbers above the rocks denote their true height in cm.




Fig. 2. Underlying surface plane fit (red) applied to 3D range data from SAD5 stereo (white point clouds) for two different altitudes.

B. Stereo-based Rock Detection

Stereo-based rock detection builds on the surface plane fit approach outlined above by (Fig. 3):

1. Fitting a plane to ground surface using the robust, least median squares algorithm.
2. Computing the standard deviation of the residuals

from the plane fit and thresholding the residuals at 1σ to identify outlier pixels that might be part of hazards (i.e. rocks).

3. Grouping the outliers by connected component extraction to identify candidate rocks, rejecting very small regions as noise.
4. Estimating rock height and position by averaging the n highest range points in each region. We typically use $n = 25$. This reduces noise in the estimates considerably.

Though very simple, this algorithm gives very good results, as described in Section IV-B.




Fig. 3. Stereo-based rock detection: 1) Robust plane fit to range data. 2) Threshold residuals at 1σ above the plane. 3) Extract connected components, rejecting small regions as noise. 4) Estimate rock height and decide which are hazards based on height.




Fig. 4. Sample results for shadow-based rock detection. 1) Input image from a data set acquired by an aircraft flying over Mars Hill in Death Valley. Sun incidence angle for this image is 45° . 2) Shadow pixel labeling. 3) Shadow ellipse fitting. 4) Rock modeling (position, diameter, and height).

C. Shadow-based Rock Detection

Shadow analysis offers the possibility of detecting rocks at much greater altitudes than is possible with stereo vision. Our algorithm has been described in detail elsewhere [2] and will only be summarized here. It has four steps (Fig. 4):

1. Acquiring the imagery.
2. Labeling shadow pixels by applying a modified Maximum Entropy Thresholding (gMET) algorithm [2] to find a threshold in the image histogram that separates shadowed from non-shadowed pixels. The modification adds a gamma-corrected version of the image to the original image before histogramming to improve the robustness of the segmentation.
3. Extracting compact shadow representations from the labeled shadow pixels. For hazard detection, it is adequate to fit a “best-ellipse” to each connected component of shadow pixels larger than 5 pixels in area.
4. Modeling rocks as circles whose centers and diame-

ters are derived from the shadow ellipses and known sun angles. Rock height is also estimated from the shadow length, sun elevation, and assumed ground plane tilt.

Section IV-C describes the performance of this algorithm for large sets of aerial images of Mars Hill in Death Valley and orbital imagery from Mars itself.

IV. HAZARD DETECTION PERFORMANCE

With current navigation technology, the best 3σ landing error ellipses for Mars are nearly circular with a diameter of ~ 20 km. Technology developments over the next decade could reduce this to < 200 m [5]; however, many hazards may still be distributed stochastically in that area. HD sensors have noise, which introduces uncertainty in the detection results that requires a probabilistic characterization. Our end goal is to give the landing site selection team a mathematical model of the probability of safe landing, given prior knowledge of the hazard distribution (typically obtained from orbit), a model of the HD sensor performance, and relevant parameters of the lander (e.g. diameter and ground clearance). This Section addresses HD sensor performance; Section V combines that with prior hazard knowledge to model the landing failure probability.

We require both a mathematical model of HD performance and experimental validation of the model. We have used three data sets so far for experimental validation:

1. “Rock wall”: stereo image pairs of a fairly natural-looking brick wall with four artificial rocks hung on it (Figs. 1-3, 5).
2. “Mars Hill”: aerial photographs of a Mars-like hill in Death Valley (Figs. 4, 6).
3. “HiRISE”: images of Mars taken from orbit with the HiRISE camera on the Mars Reconnaissance Orbiter (Fig. 12).

The rock wall data set moved stereo cameras at 10 m intervals from 10 m to 100 m from the wall and acquired at least 30 image pairs at each position. It used a stereo camera rig with a 1 m baseline, 1600x1200 8-bit/pixel imagery, and a $22^\circ \times 18^\circ$ field of view. This is the same size image as a descent camera in development for MSL. The field of view was chosen to be a good trade-off between covering enough terrain from 100 m altitude to include a safe landing site and giving adequate range resolution for hazard detection; performance evaluation results could lead to increasing or decreasing the field of view in the future. Ground truth range data was collected using a Leica Total Station. We also recorded sun azimuth and elevation angles as ground truth for shadow analysis. This data set was used to evaluate all of the algorithms in this paper.

The Mars Hill imagery was acquired in 1989 to evaluate landing hazard detectability for a Mars mission concept under study at that time. It contains 172 images on film, from which we digitized and manually registered 7 image chips

with sun incidence angles between 25° and 65° , then manually outlined 136 rocks to create ground truth for shadow-based rock detection. This data set does not have ground truth rock height or terrain slope information.

The HiRISE data set was acquired recently to support landing site selection for the Mars Phoenix mission. This camera acquires imagery from orbit (~ 300 km) at 30 cm/pixel with a line scan camera with 20,264 pixels across; typical images cover 6.2 km x 12.4 km. The data set includes 46 images covering $\sim 1,500$ km² with sun incidence angles of about 60° . Initially, geologists were manually counting rocks in very small portions of this imagery to calibrate indirect estimates of rock distributions in candidate landing sites. As discussed below, we showed that automatic shadow-based rock detection was superior, and eventually cataloged over 10 million rocks ≥ 1 m in diameter.



Fig. 5. Four frames from the rock wall dataset.




Fig. 6. Three frames from the Mars Hill data set, plus the ground truth rock masks that were extracted manually. Left to right, sun incidence angles are 25° , 40° , and 55° .

A. Slope Estimation

Ideally, performance evaluation for slope estimation should include estimator variance due to noise in the data and potential estimator bias due to incomplete filtering of rocks as outliers in the plane fit. This evaluation should be parameterized by the level of noise in the data, the number and image layout of pixels in the plane fit, and the field of view of the camera. The rock wall data set includes enough ground truth to experimentally evaluate all of this. To date, we have evaluated variance quantitatively, but we have only qualitatively verified that bias is small, so here we focus on variance.

We used linearized error propagation to model the variance of the estimated slope angle as a function of the stereo camera baseline and focal length, the subpixel precision of the stereo disparity estimates, the size of a rectangular patch in the image used to fit a plane, and the number of pixels within that patch used in the fit. The pixels were sampled with uniform spacing in the patch if not all pixels were used (i.e. for computational efficiency). Noise in the subpixel stereo disparity estimates was assumed to be zero mean Gaussian, independent from pixel to pixel, with constant variance over the entire image. This is a useful simplification of reality for an initial modeling effort.

Space does not permit including the derivation of the slope variance model, but its prediction agreed well with experimental results from the rock wall data set, which gives us confidence in the model. For each set of 30 stereo pairs at 10 m intervals from the wall, we computed the sample variance of the slope estimates. The linearized model described above has the variance of the stereo disparity as its one free parameter, so this was set to best fit the model to the experimental data. Fig. 7 plots the standard deviation of the experimental data (solid curve) against a model prediction for the area of the wall (dotted curve). The agreement is very good. A real mission might use a larger patch to estimate slope at lander scale, so Fig. 7 also plots a model prediction for a patch double the size ($\sim 6 \text{ m} \times 6 \text{ m}$). This shows a slope standard deviation $< 1.5^\circ$ at 100 m. This level of precision is sufficient to be useful for a mission, since lander slope tolerance thresholds are usually at least 15 degrees.

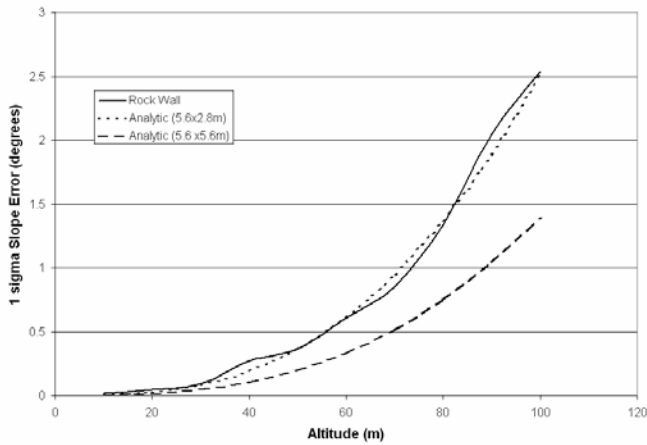


Fig. 7. Standard deviation of slope estimation error as a function of altitude, comparing of predictions of the mathematical model (dashed and dotted curves) and experimental results from the rock wall (solid curve).

This experiment used patches covering a fixed area on the ground, which means that image patches shrink with increasing altitude. This gives an error growth that is at least quadratic with altitude, as shown in Fig. 7. An alternative system architecture could fix the size of the patch in the image. The variance model predicts that this will produce slope errors that grow roughly linearly with altitude. In this case, for a patch covering one quarter of the pixels in the image, the model predicted a standard deviation of only 0.6 degrees at an altitude of 500 m, which is near the limiting altitude for this camera configuration. This would allow estimating regional slopes with good precision at high altitude, to provide an early triage of potential landing sites, then zooming in to lander-scale patches as the altitude became low enough to provide adequate precision. This architecture warrants experimental evaluation with a future data set.

B. Stereo-based Rock Detection

Rocks are hazards if their height exceeds the ground clearance or “rock tolerance” (T) of the lander. Since our approach to rock detection with stereo involves testing whether the height H of outlier regions above a fitted ground

plane exceeds a threshold $t < T$, this is a standard detection problem for which performance can be characterized by estimating the probability of detection (P_D) and the false alarm rate (FAR) as a function of t , T , altitude, sensor noise, and true rock distribution. To develop a mathematical model, we treat rock height measurements from stereo vision as having zero-mean, Gaussian noise with a standard deviation that grows quadratically with altitude. For simplicity, we ignore uncertainty in the plane fit. We then estimate P_D and FAR by taking appropriate integrals over the noisy measurement distributions. For experimental validation, we use the rock wall data set to count P_D and FAR as a function of t , T , and altitude. In ongoing work, we plan to acquire aerial stereo imagery and ground truth height data for a large set of rocks on Mars Hill as a more comprehensive test.




Fig. 8. Integrals for P_D (left) and FAR models (right).

Fig. 8 illustrates the required integrals. All rocks with true height H above T are hazards (left graph). The blue curve is the noisy measurement distribution for such a rock; the hatched area of the distribution gives P_D for that rock. This must be averaged over all rocks with true height above T to obtain P_D for the whole image:

$$P_D = \left(\sum_{i=1}^n \int_t^\infty N(x; H_i, \sigma_i) dx \right) / n \quad H_i \geq T$$

where $N(x; H_i, \sigma_i)$ is the normal distribution with mean equal to the true rock height H_i and standard deviation equal to the measurement noise σ_i at the given altitude. Rocks with true heights below T generate false alarms from the tails of their measurement distributions that exceed t (right graph). False alarms can also come from range pixels on the ground if the noise level is high enough, i.e. at relatively high altitude. We model the FAR as a combination of two terms, in order to model false alarms from small rocks and false alarms from the rest of the ground: (1) the sum over all false alarms rocks of the integrated tails of the measurement distributions for each rock, plus (2) an arbitrary number c of false alarms on the ground, where we have tuned c to fit experimental data.

$$FAR = \sum_{j=1}^m \int_t^\infty N(x; H_j, \sigma_j) dx + c \int_t^\infty N(x; 0, \sigma) dx \quad H_j < T$$

The rock wall data set has rocks 14, 25, 37, and 54 cm high. Fig. 1 showed a 3-D rendering of the range data from 40 m; Fig 9 shows a rendering from 90 m. Since lander rock

tolerances in recent missions have been around 50 cm, we treated the 54 cm rock as a hazard and measured P_D and FAR as a function of distance from the wall, averaging over all 30 frames at each distance. Fig. 10 plots these for $T = 50$ cm and $t = 45$ cm, together with a corresponding model prediction. The hazardous rock was detected in every image, with no false alarms, up to 70 m away. The model and the experimental results deviate from each other appreciably by 90 m; this may be due to subtle biases in rock height estimates that become more significant at higher altitudes. A key question is how these results would change with rock height distributions in real Mars terrain; we address this in Section V.




Fig. 9. 3-D rendering from below of stereo range data of the rock wall from 90 m distance. 54 cm and 37 cm rocks are visible.

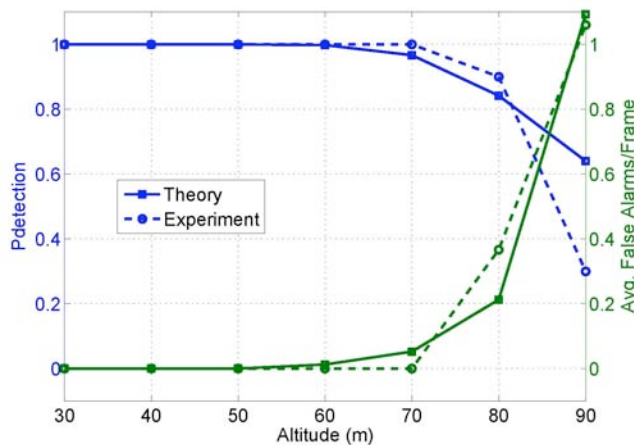


Fig. 10. Hazard detection and false alarm rate comparison between the model and experimental results for the wall data set.

C. Shadow-Based Rock Detection

The reliability of shadow-based rock detection depends on many factors, including several that affect detectability of the shadows and several that affect accuracy of height estimates made from the shadows. A full evaluation of these factors is beyond the scope of this paper, so we focus here on two key issues that are testable with our data sets: height estimation accuracy with the rock wall data and shadow detection reliability with the Mars Hill and HiRISE data.

With the rock wall data, all of the shadows were reliably detected without false alarms at all ranges. Over all four rocks, the RMS height estimation error was 1.28 cm at 50 m, or 3.9% of the average rock height of 33 cm, and 1.65 cm at 100 m, or 5% of the average rock height. This is encouraging, but it needs to be tested with a variety of sun angles and in real terrain with a larger number of rocks.

Fig. 11 shows shadow detection results for the Mars Hill data as a function of rock diameter in units of pixels (with 5

pixel bin widths), averaged over all 7 sun incidence angles. P_D over the entire data set was 85%, with a FAR of 3.26 per image. However, 56% of the missed detections occurred at 5-10 pixels in diameter. The largest fraction of those occurred for the image with the longest shadows (65° incidence), where the shadow of the largest rock obscured many nearby small rocks. Mission design constraints due to Earth-Mars communication considerations likely will lead to landing at times that give shorter shadows, with sun incidence around 45° . The vast majority of the false positives gave rock diameter estimates ≤ 15 pixels. For typical rocks, typical landers, and the angular resolution we expect in a descent camera, hazardous rocks will be around 40 pixels in diameter from 100 m altitude. In summary, the errors in this experiment would be unlikely to be important in real missions.

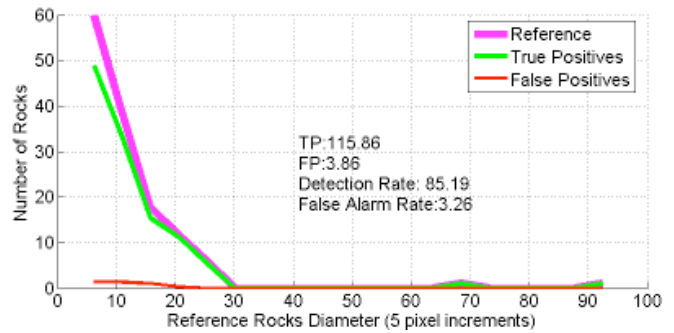


Fig. 11. Shadow-based rock detection. Average detection for seven different sun angles.

With the HiRISE imagery, on the order of 4000 rocks were counted by hand by geologists to compare to automatic, shadow-based rock detection. 1226 of the largest rocks were used to determine that rock diameters estimated manually (by looking at the rocks themselves) agreed well with diameters estimated automatically from the shadows. Manually and automatically estimated distributions of diameters of the full set were compared and found to match very closely, which gives a statistical confirmation of high quality automatic detection. Fig. 12 shows a sample result.




Fig. 12. Sample result of shadow detection and rock estimation for a 323x322 pixel HiRISE sub-image. The largest shadow is about 9 pixels across. On the right, black spots are detected shadows and cyan spots are estimated rocks for all shadow regions > 5 pixels in area.

V. SAFE LANDING PROBABILITY MODEL

Given a performance model for the HD sensor, a model of the hazard size/frequency distribution, and lander param-

ters, predictions of the probability of safe landing can be made either by mathematical modeling or by Monte Carlo simulation. Here we develop a mathematical model; future work may compare this to simulation. We idealize the process of selecting a landing site onboard by dividing the descent image into m non-overlapping square regions, each big enough to land in. The probability P_S of landing safely is the probability P_{ND} of finding one region where no hazards are detected times the probability P_{NMD} that the selected region has no missed detections: $P_s = P_{ND} P_{NMD}$. From Section IV-B, we have a model of rock detection with stereo vision that is acceptable up to 70 m altitude. Models of Mars rock populations have been developed by geologists from theoretical considerations and from rocks seen in Mars surface imagery [9]. We will now use those models to derive expressions for P_{ND} and P_{NMD} for the case of rock detection with stereo vision. We again assume zero mean, Gaussian noise in rock height estimates and use the notation t , T , and H for a given rock of height H , the probability of experiencing a false positive (P_{FP}) or false negative (P_{FN}) follows similar logic to section IV-B and is given by:

$$P_{FN}(H,t) = \int_0^t N(x;H,\sigma) dx \quad \text{for } H \geq T$$

$$P_{FP}(H,t) = \int_t^\infty N(x;H,\sigma) dx \quad \text{for } H < T$$

Weighting P_{FN} and P_{FP} by the frequency of rocks of height H per unit area and integrating over H would give us the expected number N_{FN} of false negatives and the expected number N_{FP} of false positives per unit area. Planetary geologists characterize rock distributions by the cumulative fractional surface area F covered by rocks of diameter $\geq D$; that is, for diameter D , by the fraction of the surface area covered by rocks of diameter D or larger. It is well established [9] that F is well modeled by

$$F(k,D) = k \cdot \exp(-q(k) \cdot D)$$

with

$$q(k) = c_a + c_b / k$$

where k is the *rock abundance*, the percent of the surface area covered by rocks of all diameters, and c_a and c_b are other constants used to fit the data. Different landing sites are characterized by their value of k . For the Mars Viking 1 (VL1), Viking 2 (VL2), and Pathfinder (MPF) landing sites, estimates of k are 17%, 18%, and 20%, respectively; close to the Bonneville Crater that was observed by the Spirit rover, k was 32%. Mars Hill is about 22%. These are all very high compared to typical Mars values estimated from orbital thermal inertia measurements, which put the modal value at 6% and the median around 8%. Other places in the solar system of scientific interest have been observed to have rock distributions similar to MPF, such as asteroid Eros and the lunar highlands.

For our purposes, we convert F to a frequency distribution f giving the number of rocks of diameter D per square meter:

$$f(k,D) = \frac{d(1-F(k,D))}{dD} / (0.25\pi D^2)$$

$$= k \cdot q(k) \cdot \exp(-q(k) \cdot D) / (0.25\pi D^2)$$

From observations around the Mars landing sites, rock

height on average is half the diameter, so we make the approximation that $D = 2H$ and convert f to a function of H by simple substitution. We now have:

$$N_{FN}(k,t,T) = \int_T^\infty P_{FN}(H,t) \cdot f(k,H) dH$$

$$N_{FP}(k,t,T) = \int_0^T P_{FP}(H,t) \cdot f(k,H) dH$$

The expected number of hazardous rocks per unit area is:

$$N_{HZ}(k,T) = \int_T^\infty f(k,H) dH$$

The expected number of detected hazardous rocks per unit area is then:

$$N_{DH}(k,t,T) = N_{HZ}(k,T) + N_{FP}(k,t,T) - N_{FN}(k,t,T)$$

From [13], the number of hazardous rocks in a given area A can be modeled by a Poisson distribution with mean AN_{HZ} , from which it follows that the probability of exactly n hazardous rocks in an area A is given by:

$$P(n,\lambda) = \frac{\lambda^n \cdot \exp(-\lambda)}{n!} \quad \text{with } \lambda = AN_{HZ}$$

and the probability of at least one hazardous rock in area A is

$$1 - P(0,\lambda) = 1 - \exp(-AN_{HZ})$$

If A is the area of the lander, the above expression is the probability that a blind landing will fail due to rock impalement. Denote this expression as P_f and note that it has parameters k and T , since these are parameters of N_{HZ} .

Recall that our idealization for landing with an HD sensor is to divide the field of view of the sensor into m non-overlapping potential landing sites. From [13] if hazard detection was perfect, the landing would fail if all m sites had at least one hazardous rock, which has probability

$$(P_f)^m$$

Landing successfully means finding at least one site free from rock hazards, which is simply

$$1 - (P_f)^m$$

For noisy rock detection with stereo vision, P_f gets replaced by

$$1 - \exp(-AN_{DH})$$

as the probability of detecting a hazard in any of the m candidate landing sites, so the probability of finding one of m sites with no detections is

$$P_{ND} = 1 - (1 - \exp(-AN_{DH}))^m$$

The probability of no missed detections in any given candidate landing site the Poisson distribution $P(0, N_{FN})$, so

$$P_{NMD} = \exp(-AN_{FN})$$

With the probability of safe landing as $P_s = P_{ND} P_{NMD}$, this completes our model.

Fig. 13 shows plots that instantiate this model for parameters of the MSL mission, which are a rock tolerance $T = 60$ cm, a lander undercarriage area of 4 m^2 , and a rock height hazard threshold of 55 cm. The MSL goal of limiting failure probability due to rock impalement to $< 0.25\%$ (i.e. safe landing probability of > 0.9975) constrains blind landers to a rock abundance of about 6%, which is about 1/3 of the planet. The stereo HD plot (green) is for a sensing altitude of 70 m and exceeds the MSL goal for all rock abundances shown. Note that this prediction is consistent with the rock

wall results at 70 m observed in Section IV-B. The reduction in failure probability for landing with HD compared to blind landing is easily more than a factor of four in this model for rock abundances above 6%. If rocks were the only hazards on Mars, the model implies that using this HD capability would increase the area of Mars accessible to safe landing to nearly 100%. We have not yet assessed what area would be eliminated by other hazards, particularly slopes.

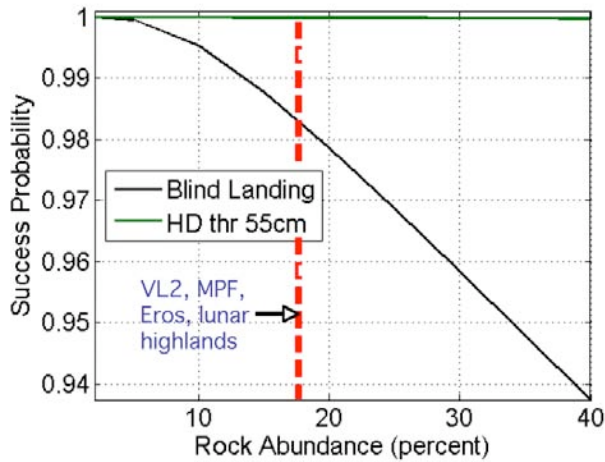


Fig. 13. Probability of safe landing as a function of rock abundance, comparing landing with HD capability to blind landing, instantiated for MSL parameters (60 cm rock tolerance, 4 m² lander undercarriage, 55 cm rock height threshold). The green plot predicts stereo-based probability with sensing from 70 m altitude. The black plot is for blind landing.

VI. CONCLUSION

We used a Mars landing scenario as an extreme case of a fast, near-vertical descent to motivate sensor selection for landing hazard detection. This and considerations of minimizing mass, power, and volume while maximizing relevance to other missions led us to conclude that stereo vision and shadow analysis with descent cameras appear to be the smallest sensor suite with the widest applicability, given the state of development of sensor alternatives today. We then outlined algorithms we have developed to date to detect slope hazards with stereo vision and rock hazards with stereo vision and shadow analysis. We derived analytical performance models for these based on Gaussian noise models, compared the prediction of those models to experimental data, and found reasonably good agreement. This implies that the models are useful for predicting performance of these functions in operational scenarios. We then embedded the hazard detection performance models in a model for the probability of landing safely, given parameterized models of lander rock tolerance, lander area, and parameterized rock size/frequency distributions fit to Mars and terrestrial data. When this model is instantiated for parameters of the MSL mission, it predicts that even very conservative assumptions about the performance of the vision system will reduce the probability of a failed landing by at least a factor of four compared to a blind landing for any rock abundance. Conversely, for the level of safety desired by MSL, it predicts

that the vision system would allow access to nearly the entire planet, compared to about one-third of the planet blind landing. This would represent a major improvement in access to sites of scientific value for a small increase in sensor payload. Analogous benefits should accrue to missions to other bodies in the solar system.

REFERENCES

- [1] A. E. Johnson, A. Klumpp, J. Collier, A. Wolf, "Lidar-based Hazard Avoidance for Safe Landing on Mars", *AIAA Journal of Guidance, Control and Dynamics*, 25(5), October 2002.
- [2] A. Huertas, Y. Cheng, R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *Proc. IEEE Aerospace Conference*, Big Sky, MO. Mar. 2006.
- [3] J. De Lafontaine, A. Ulitsky, J.W. Tripp, R. Richards, M. Dal., C. Sallaberger, "LAPS: The Development of a Scanning LIDAR System with GNC for Autonomous Hazard Avoidance and Precision Landing", *Proc. SPIE Vol. 5418, Spaceborne Sensors*, 2004, pp. 81-93.
- [4] B. D. Pollard, G. Sadowy, D. Moller, E. Rodriguez, "A Millimeter-Wave Phased Array Radar for Hazard Detection and Avoidance on Planetary Landers", *IEEE Aerospace Conference*, March 2005, pp. 1115-1122.
- [5] S. W. Thurman, L. H. Matthies, "Space Flight Test of Vision-Guided Planetary Landing System," *Proc. Infotech@Aerospace2007 Conference*, Rohnert Park, CA, May 2007.
- [6] M. C. Malin, et al., "The Mast Cameras and Mars Descent Imager (MARDI) for the 2009 Mars Science Laboratory", *Lunar and Planetary Science XXXVI*, Extended Abstract N. 1214, Lunar and Planetary Institute, Houston, TX, 2005.
- [7] J. Veverka et al. "Imaging of Small-Scale Features on 433 Eros from NEAR: Evidence for a Complex Regolith," *Science*, Vol. 292. no. 5516, Apr. 2001, pp. 484 - 488.
- [8] P. Christensen, et al., "Morphology and Composition of the Surface of Mars: Mars Odyssey THEMIS Results", *Science*, 300(2056), 2003.
- [9] M. Golombek, et al., "Rock Size-Frequency Distributions on Mars and Implications on Mars Exploration Rover Landing Safety and Operations," *Journal of Geophysical Research*, 108(E12), 2003.
- [10] L. Matthies, M. Maimone, A. Johnson, Y. Cheng, R. Willson, C. Villalpando, S. Goldberg, A. Huertas, "Computer Vision on Mars", *Int'l Journal of Computer Vision*, 75(1), pp. 67-92, October 2007.
- [11] H. Hirschmuller, P. Innocent, J. Garibaldi, "Real-time correlation-based stereo vision with reduced border errors," *Int'l Journal of Computer Vision*, 47(1-3), pp. 229-246, April-June 2002.
- [12] C. Villalpando, "Acceleration of Stereo Correlation in Verilog", *9th Military and Aerospace Programmable Logic Devices (MAPLD) International Conference*, Washington, D.C., September 2006.
- [13] D. Bernard, M. Golombek, "Crater and Rock Hazard Modeling For Mars Landing," *AIAA Space Conference*, Albuquerque, NM Aug. 2001.