

# Evidential Terrain Safety Assessment for an Autonomous Planetary Lander

Navid Serrano, Anthony Quivers\*, Max Bajracharya, Homayoun Seraji  
*Jet Propulsion Laboratory*  
*California Institute of Technology*  
*4800 Oak Grove Dr*  
*Pasadena, CA 91109, USA*  
*firstname.lastname@jpl.nasa.gov*

*\*Department of Aeronautics and Astronautics*  
*Massachusetts Institute of Technology*  
*77 Massachusetts Ave*  
*Cambridge, MA 02139, USA*  
*quiversa@mit.edu*

## Abstract

*In this paper, the problem of determining safe regions of the visible terrain for autonomous spacecraft landing is explored. A multi-sensor architecture is proposed with a combination of active and passive sensors that monitor the terrain during descent. Craters and rocks are identified using hazard detection algorithms that provide a preliminary assessment of unsafe regions of the terrain. Slope and roughness features are also extracted and provide a richer description of the surface topography. However, a direct relationship between the topographic features and terrain safety is unknown. Hence, supervised learning techniques are employed in order to obtain an optimal mapping between the topographic features and terrain safety. An evidential framework is used to fuse the individual terrain safety assessments yielded by each sensor using Dempster-Shafer theory. The evidential terrain safety assessment provides a continuous-valued landing score for all visible points on the terrain that can be used to determine the safest landing site.*

## 1. Introduction

The Entry, Descent, and Landing (EDL) phase is critical for the success of a landed space exploration mission. Recent missions have generally opted for a “hard” landing, including both the Pathfinder and Mars Exploration Rovers (MER) missions. For future missions, NASA has focused on EDL technologies that enable “soft” landing driven by autonomous terrain characterization capabilities.

The work presented here is part of an on-going effort to investigate the use of multiple on-board sensors for autonomous hazard assessment and safe site

selection [1-3]. Sensor fusion was discussed in [1], and Bayesian and Fuzzy approaches to landing site selection were presented in [2] and [3], respectively. In this paper, the use of evidential reasoning for landing safety assessment is investigated. Landing safety assessment is the process of determining how safe the visible terrain is for landing. The terrain is characterized in real-time during descent using a suite of on-board sensors. Points on the visible terrain must be assigned a safety score in order to determine a viable landing target.

The multi-sensor architecture provides a rich set of measurements from which to characterize the terrain. A combination of active (radar and lidar) and passive (camera) sensors are used. Hazard detection algorithms are used to determine the presence of craters and rocks on the terrain from visible imagery. The absence of these hazards, however, does not necessarily indicate that the terrain is certainly safe. Hence, topographic features (slope and roughness) from range data are also extracted to enhance the terrain safety assessment. However, the relationship between the topographic features and terrain safety is not explicitly known. Thus, supervised learning techniques are used to establish the best mapping between the topographic features and terrain safety. Decision trees [4] are considered because their computational simplicity makes them suitable for on-board computers. The safety ground truth is obtained using a simple geometric algorithm that estimates the final pose of the spacecraft based on the layout of the terrain.

The individual contributions from each sensor are fused using Dempster-Shafer (DS) theory [5] in order to arrive at a final safety score. DS theory provides a framework for evidential reasoning that explicitly accounts for both known and unknown information.

## 2. On-board sensors

As intimated earlier, multiple on-board sensors are considered in order to provide both richness and robustness. The specifications of the sensors are shown in Table 1.

**Table 1. On-board sensor characteristics**

Sensor	Type	Max. Range	FOV	Image Size
Radar	Active	10km	27°	10×10
Camera	Passive	8km	37°	1024×1024
Lidar	Active	1km	12°	100×100

Based on the inherent range of operation of each sensor shown in Table 1, the descent is divided into three tiers. In tier 1, roughly between 10km and 8km, only the radar is operational. At 8km the camera becomes operational and together with the radar constitutes tier 2. Finally, in tier 3, from 1km to touchdown, all three sensors are operational, as summarized in Table 2.

**Table 2. Tiered sensor operation**

Tier	Range	Operational Sensor(s)
1	10km – 8km	Radar
2	8km – 1km	Radar + Camera
3	1km - Touchdown	Radar + Camera + Lidar

During spacecraft descent, there are only a limited number of retargeting maneuvers that can be performed. Using the tiered approach described here, three retargeting maneuvers could be allotted—one per tier.

## 3. Hazard detection

### 3.1. Crater detection

The crater detection algorithm is described in [6]. Given the known location of the sun during descent, the algorithm can accurately locate craters based on their shadow patterns. Each detected crater is modeled as an ellipse. Let  $(x_i, y_i)$ ,  $a_i$ ,  $b_i$ , and  $\phi_i$  be the ellipse center, semi-major axis length, semi-minor axis length, and rotation angle, respectively, for the  $i$ th detected crater. The crater map is defined as:

$$f_c(x, y) = \begin{cases} \beta_i, & \left( \frac{x^2}{a_i^2} + \frac{y^2}{b_i^2} \right) \leq 1, \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where  $x$  and  $y$  are points in a coordinate system rotated by  $\phi_i$  and translated by  $x_i$  and  $y_i$ . Non-zero values of  $f_c$  indicate the presence of craters with a degree of certainty  $\beta_i \in [0,1]$ .

### 3.2. Rock detection

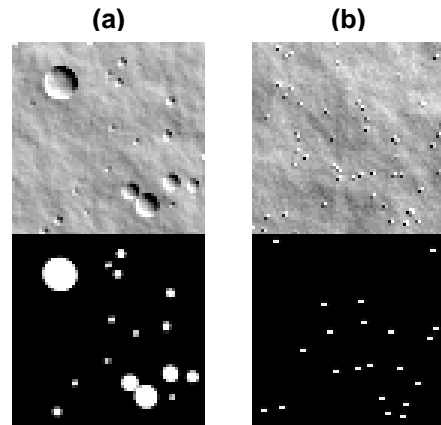
The algorithm described in [7] uses segmentation to first identify shadow regions in the image. Given the known sun angle, the location and size of rocks are determined based on shadow projection patterns. Let  $b(x,y)$  represent the binary map of detected rocks produced by the algorithm:

$$b(x, y) = \begin{cases} 1, & (x, y) \in R \\ 0, & \text{otherwise} \end{cases}, \quad (2)$$

where  $R$  is the set of pixel locations identified as rocks. The binary map  $b(x,y)$  is filtered to account for uncertainty in the detection process and err on the side of caution. The final rock map is:

$$f_r(x, y) = \sum_i \sum_j b(i, j) h(x-i, y-j), \quad (3)$$

where  $h(x,y)$  is any suitable low-pass filter whose elements sum to unity, which ensures that  $f_r(x,y) \in [0,1]$ . Crater and rock detection results are shown in Figure 1.



**Figure 1. Example crater (a) and rock (b) detection results.**

Hazards detected from camera imagery provide a preliminary characterization of the terrain. Yet, it should be borne in mind that the crater and rock maps  $f_c$  and  $f_r$  only indicate the location of *unsafe* landing hazards. If a hazard (i.e. crater or rock) is not detected, the terrain cannot be assumed to be necessarily safe. For a complete terrain safety assessment, further information is necessary. The radar and lidar range information is used for this purpose.

## 4. Terrain classification

The process of classifying points on the terrain as either safe or unsafe based on simple topographic features derived from range data is discussed in this section.

### 4.1. Terrain features

The topographic features must be highly descriptive yet simple enough to satisfy on-board computational constraints. The key indicator of safety is the slope of the terrain. The slope of the terrain can be estimated from range data using plane fitting techniques. The local terrain slope is defined as the angle formed by the estimated plane normal  $\mathbf{n}$  and geodetic normal  $\mathbf{n}_g$ :

$$f_\theta(x, y) = \cos^{-1}(\mathbf{n}^T \mathbf{n}_g). \quad (4)$$

As shown in [8], the Least Median of Squares (LMedSq) regression technique approximates the orientation of the terrain underneath objects on the terrain surface (such as rocks). Thus, a terrain roughness feature can be defined based on the residual between the sensor range value  $d(x, y)$  and the value of the fitted plane using LMedSq regression:


$$f_e(x, y) = |d(x, y) - \hat{d}(x, y)|, \quad (5)$$

where  $\hat{d}(x, y)$  is the value of the fitted plane at location  $(x, y)$ . The slope is extracted from radar range data and the roughness is extracted from lidar range data.

### 4.2. Safety ground truth

The mapping between the aforementioned topographic features and terrain safety is learned in a supervised fashion. In order to learn the mapping between the features and safety, a means of obtaining safety ground truth is needed. For this purpose, a

simple geometric algorithm to determine the lander's final pose on the terrain surface is proposed. The lander is modeled as a plane and determine three supporting contact points.



**Figure 2. Determination of spacecraft contact points.**

The first contact point  $\mathbf{p}_1$  is the highest point on the terrain within the boundary of the plane defining the lander's footprint. Having touched  $\mathbf{p}_1$ , the spacecraft will tilt in the direction of the center of gravity, as shown in Figure 2a. A new search space is defined by the points  $\mathbf{p}_i$  that satisfy

$$(\mathbf{p}_i - \mathbf{p}_1)^T (\mathbf{p}_{cg} - \mathbf{p}_1) > 0, \quad (6)$$

where  $\mathbf{p}_{cg}$  is the location of the center of gravity. The second contact point  $\mathbf{p}_2$  is the one with maximum tilt:

$$\phi_i = \tan^{-1} \left( \frac{\|\mathbf{p}_i - \mathbf{p}_1\|}{z_1 - z_i} \right), \quad (7)$$

where  $z_1$  and  $z_i$  are the terrain elevation values of the first and  $i$ th points, respectively. Having determined the second contact point  $\mathbf{p}_2$ , as shown in Figure 2b, the new search space includes the points  $\mathbf{p}_i$  that satisfy

$$\left[ (\mathbf{p}_2 - \mathbf{p}_1)^T (\mathbf{p}_i - \mathbf{p}_1) \right] \left[ (\mathbf{p}_2 - \mathbf{p}_1)^T (\mathbf{p}_1 - \mathbf{p}_{cg}) \right] > 0. \quad (8)$$

As before, the maximum tilt is used to determine the third contact point  $\mathbf{p}_3$ . But in this case, the tilt is now

$$\phi_i = \tan^{-1} \left( \frac{\mathbf{p}_{12}^T \mathbf{p}_{12} \mathbf{p}_{21}^T \mathbf{p}_{li}}{(z_2 - z_1) \mathbf{p}_{21}^T \mathbf{p}_{li} + (z_1 - z_i) \mathbf{p}_{12}^T \mathbf{p}_{12}} \right), \quad (9)$$


where  $\mathbf{p}_{ij} \equiv \mathbf{p}_i - \mathbf{p}_j$ . The landing orientation is determined based on the normal to the plane that intersects the three contact points shown in Figure 2c:

$$\mathbf{n} = (\mathbf{p}_3 - \mathbf{p}_1) \times (\mathbf{p}_2 - \mathbf{p}_1). \quad (10)$$

Terrain safety is determined based on the stability of the lander. For this analysis, the lander is modeled as a cuboid in  $\mathcal{R}^3$ . The center of gravity lies at a distance  $w$  from the bottom plane. The lander is stable (and the terrain is safe) if the center of gravity lies within the boundary of the bottom face of the cuboid (as in Figure 3); otherwise it will topple (and the terrain is unsafe). In order to make this determination, the location of the center of gravity is projected onto the  $xy$ -plane:

$$\mathbf{p}'_{cg} = wP \frac{\mathbf{n}}{\|\mathbf{n}\|}, \quad (11)$$

where  $P$  is an orthographic projection matrix and  $\mathbf{n}$  is the plane normal (10). Therefore, the terrain is safe for landing if  $\mathbf{p}'_{cg}$  lies within the boundary defined by bottom face of the cuboid, as shown in Figure 3.



**Figure 3. Terrain safety based on the lander's final orientation on surface.**

#### 4.3. Terrain classification using decision trees

Numerous classification engines exist that could be used to learn a mapping from terrain features to terrain safety that minimizes classification error. Well known classifiers include Neural Networks [9] and kernel methods, such as Support Vector Machines [9]. However, these approaches are better suited to more complex classification problems, particularly non-linear mappings in very high-dimensional spaces. Furthermore, even the most efficient implementation of any of these classification engines still requires computational resources that may be beyond those available on a mission spacecraft.

For the application considered here, a very small set of features is used and the mapping between these features and terrain safety is relatively simple. It is intuitive, for instance, that unsafe regions of the terrain should map to high slope values. What is not known, however, is exactly what slope value will minimize the error. It is in this respect that decision trees are very useful.

A decision tree is a set of logical conditions that can be used to map a set of predictors to a categorical target. Decision trees effectively partition the feature space into multiple regions, each associated with one of the target values, or classes. Decision trees are particularly attractive because of their simplicity. Inference merely involves evaluating a sequence of if-then statements. Each if-then statement represents a partition of the feature space, which is effectively a threshold.

Many methods exist to train decision trees. One such approach is known as Classification and Regression Trees (CART) [4]. CART yields an optimal decision tree that best balances classification error and complexity (i.e. tree size). CART uses a set of statistical measures to determine optimal partitions of the feature space. A decision tree can grow arbitrarily large until every observation in the training set is properly classified—an impractical scenario. Therefore, CART employs stopping criteria based on classification error tolerance. Furthermore, the decision tree can be pruned in order to arrive at a tree size that suits the task at hand, while minimizing classification error.

For the task of terrain classification, CART is used to learn decision trees that map radar and lidar features to terrain safety. This is similar to the approach used in [8] to map lidar slope and roughness features to terrain safety; except in this case the thresholds are optimally determined using CART. Application of the decision trees will yield a binary safety map for both the radar and lidar. (The camera safety map is obtained directly from the hazard detection algorithms.) Since each sensor has its own safety map, a means of fusing the information is needed.

## 5. Evidential safety assessment

### 5.1. Dempster-Shafer sensor fusion

In DS theory [5], evidence from a sensor is measured as belief. Each sensor contributes a belief of 1.0 and distributes this belief among certain propositions. The combined mass of belief for two sensors is obtained using Dempster's combination rule:

$$m^{ij}(A) = \frac{\sum_{B \cap C = A} m^i(B)m^j(C)}{1 - \sum_{B \cap C = \emptyset} m^i(B)m^j(C)}, \quad (12)$$

where  $m(A)$  represents the proportion of evidence that supports the claim that a particular element of the universal set  $X$  belongs to the set  $A$ , and  $i$  and  $j$  represent two arbitrary sensors. In our case, the universal set is  $X = \{safe, unsafe, unknown\}$ . Beliefs need not be mutually exclusive; thus,  $m(unknown) = 1 - m(safe) - m(unsafe)$ .

The belief masses for each sensor are obtained from the terrain features discussed above. The radar belief mass is:

$$m^R(safe) = \begin{cases} 1, & f_\theta \leq T^R \\ 0, & otherwise \end{cases}, \quad (13)$$

where  $f_\theta$  is the radar slope defined earlier and  $T^R$  is a set of thresholds represented by nodes in the decision tree yielded by CART. The *unsafe* belief mass is defined as  $m^R(unsafe) = 1 - m^R(safe)$ . Whenever radar measurements are available,  $m^R(unknown) = 0$ . Conversely, for points lying outside the radar's field of view,  $m^R(unknown) = 1$ . The camera belief mass is:

$$m^C(unsafe) = \max\{f_c, f_r\}, \quad (14)$$

where  $f_c$  and  $f_r$  are the crater and rock detection maps. Only an *unsafe* belief mass is defined for the camera because, as discussed earlier, the hazard detection algorithms do not provide information about *safe* regions. Hence,  $m^C(unknown) = 1 - m^C(unsafe)$  and  $m^C(safe) = 0$ . Finally, the lidar belief mass is:

$$m^L(safe) = \begin{cases} 1, & f_e \leq T^L \\ 0, & otherwise \end{cases}, \quad (15)$$

where  $f_e$  is the lidar roughness defined earlier and  $T^L$  is the set of thresholds learned by CART for the lidar. Similarly to the radar, the *unsafe* belief mass is defined as  $m^L(unsafe) = 1 - m^L(safe)$ . Again,  $m^L(unknown) = 0$  for points lying within lidar's field of view, and  $m^L(unknown) = 1$  otherwise.

## 5.2. Evidential safety score

Dempster's combination rule (12) can be applied iteratively for more than two sensors. Once the combined mass of belief is obtained, a safety score that

takes into account the available evidence is computed. A completely *unknown* area is considered neither *safe* nor *unsafe*, and consequently it is assigned a safety score of 0.5. If a landing site is certainly *safe* or certainly *unsafe*, it is assigned a score of 1.0 or 0.0, respectively. These scores define the boundaries. Points in between are defined linearly. The evidential safety score is shown in Figure 4 and is defined as:

$$s(x, y) = \frac{1}{2} [m^{RCL}(safe) - m^{RCL}(unsafe) + 1], \quad (16)$$

where  $m^{RCL}(safe)$  and  $m^{RCL}(unsafe)$  are obtained by iteratively applying equation (12).




Figure 4. Surface defined by evidential safety score.

## 6. Spacecraft descent simulation

### 6.1. DSEDS

In order to simulate spacecraft descent onto a planetary surface—including the spacecraft pose and sensor data—the high-fidelity Dynamics Simulator for Entry, Descent and Surface landing (DSEDS) [10] is used. DSEDS was developed at JPL and incorporates spacecraft, atmosphere, and planetary models as well as realistic guidance and navigation algorithms and has been used to simulate Mars EDL for various projects and missions. External sensor models for the radar, camera, and lidar are integrated to extract terrain features during descent.

## 6.2. Planetary terrains

A recursive fractal method [11] is used to create a base terrain Digital Elevation Model (DEM),  $t(x,y)$ . The synthetic terrains are consistent with Bernard and Golombek's definition of "moderately cratered" regions of Mars [12]. The average number of craters per unit area is used as the rate parameter of a Poisson random variable that yields the total number of craters in the DEM. These craters are then sized based on the distribution of crater diameters reported in [12] and are randomly placed on the DEM (according to a uniform distribution). The crater model is a hemisphere with a radius  $R$  and a depth of  $0.4R$  [12]. An ejecta blanket extends from the outer rim with an elevation profile proportional to  $1/r^3$ , where  $r$  is the distance from the crater center. The  $i$ th crater in the terrain is thus modeled as:

$$c_i(x,y) = \begin{cases} -d \sqrt{1 - \frac{(x^2 + y^2)}{R_i^2}} + 0.08R_i, & x^2 + y^2 \leq R_i^2 \\ \frac{0.08R_i^4}{\sqrt{(x^2 + y^2)}^3}, & x^2 + y^2 > R_i^2 \end{cases}. \quad (17)$$

The crater is displaced by 0.08 above the terrain surface in order to properly model the elevated rim and ejecta blanket.

Rock abundance is also determined using a Poisson model. However, different rate parameters are used based on proximity to craters [12]. Golombek's exponential rock model is used to randomly set the diameter of each rock [12]. The rock model is a hemisphere perturbed by noise:

$$r_i(x,y) = \frac{D_i}{2} \sqrt{1 - \frac{((x+\eta_x)^2 + (y+\eta_y)^2)}{(D_i/2)^2}}, \quad (18)$$

where  $D_i$  is the diameter of the  $i$ th rock, and  $\eta_x$  and  $\eta_y$  are normally distributed random variables that determine the roughness of the rock. The final terrain DEM is a linear combination of the fractal terrain and the craters and rocks:

$$z(x,y) = t(x,y) + \sum_i c_i(x,y) + \sum_i r_i(x,y). \quad (19)$$

## 6.3. Results

Spacecraft descent is simulated for ten different terrain DEMs using DSEDS. During a simulated

descent, the sensors are activated and terrain features are extracted. The features are all co-registered. For supervised learning, the ten DEMs are partitioned into training and testing sets (five DEMs in each). Each point on the terrain is treated independently and used as a sample for training.

**Table 3. Results using CART**

Tier	Radar	Lidar
1	93.3% (2 nodes)	N/A
2	92.4% (2 nodes)	N/A
3	88.6% (2 nodes)	90.5% (2 nodes)

The results shown in Table 3 represent the percentage of sample points (drawn from the terrains in the test set) that are correctly classified. Table 3 also shows the number of nodes in the optimal decision tree yielded by CART. The resulting number of nodes in each optimal tree is only two. This means that the mapping between the relevant feature and terrain safety involves only one partition or threshold—the simplest possible scenario. For comparison, results using a Bayesian classifier are also shown in Table 4. As can be seen, the results in Tables 3 and 4 are very close.

**Table 4. Results using Bayesian classification**

Tier	Radar	Lidar
1	93.7%	N/A
2	96.0%	N/A
3	87.4%	91.1%

Full results using evidential reasoning are shown in Figure 5. The evidential safety score obtained from operational sensors in each tier is overlaid on two different terrain DEMs. The safety score is shown at altitudes of 8000m (tier 1), 4000m (tier 2) and 500m (tier 3). The color scheme used is consistent with the surface shown in Figure 4, where green and red correspond to the two extremes *safe* (1.0) and *unsafe* (0.0), respectively and *unknown* (0.5) is shown in yellow.

As can be seen in Figure 5, the evidential safety score is highest when the individual contributions from each sensor are in agreement. Disagreement between the sensors leads to uncertainty and, in turn, the safety score tends towards the midpoint (0.5). This is evident in the outer portion of the results from tier 2 and 3 in Figure 5. This outer portion is the region where only the camera information is available and there is no contribution from the radar or lidar. Unless a crater or

rock is detected by the camera, the terrain will be deemed *unknown* and the safety score (as shown) is consequently 0.5. When a hazard is detected by the camera, the safety score appropriately tends to 0.0. This can be seen in the results for tier 3 where the terrain is mostly *unknown* (shown in yellow) unless a rock is detected (shown in red).

During descent, the spacecraft can be retargeted away from points with a low score (0.0) and towards points with a high score (1.0). Because the evidential safety score is more assertive when sensors agree and less assertive when they disagree, the approach can be regarded as aggressive for the purposes of landing site selection. That is, fewer points on the terrain are deemed *unsafe* (0.0), and consequently, there are more potential landing sites to choose from. In practice, it may be desirable to have both aggressive and conservative estimates in order to make different types of retargeting decisions based on the lander's distance to the surface. This idea will be explored in future research.

## 7. Conclusions

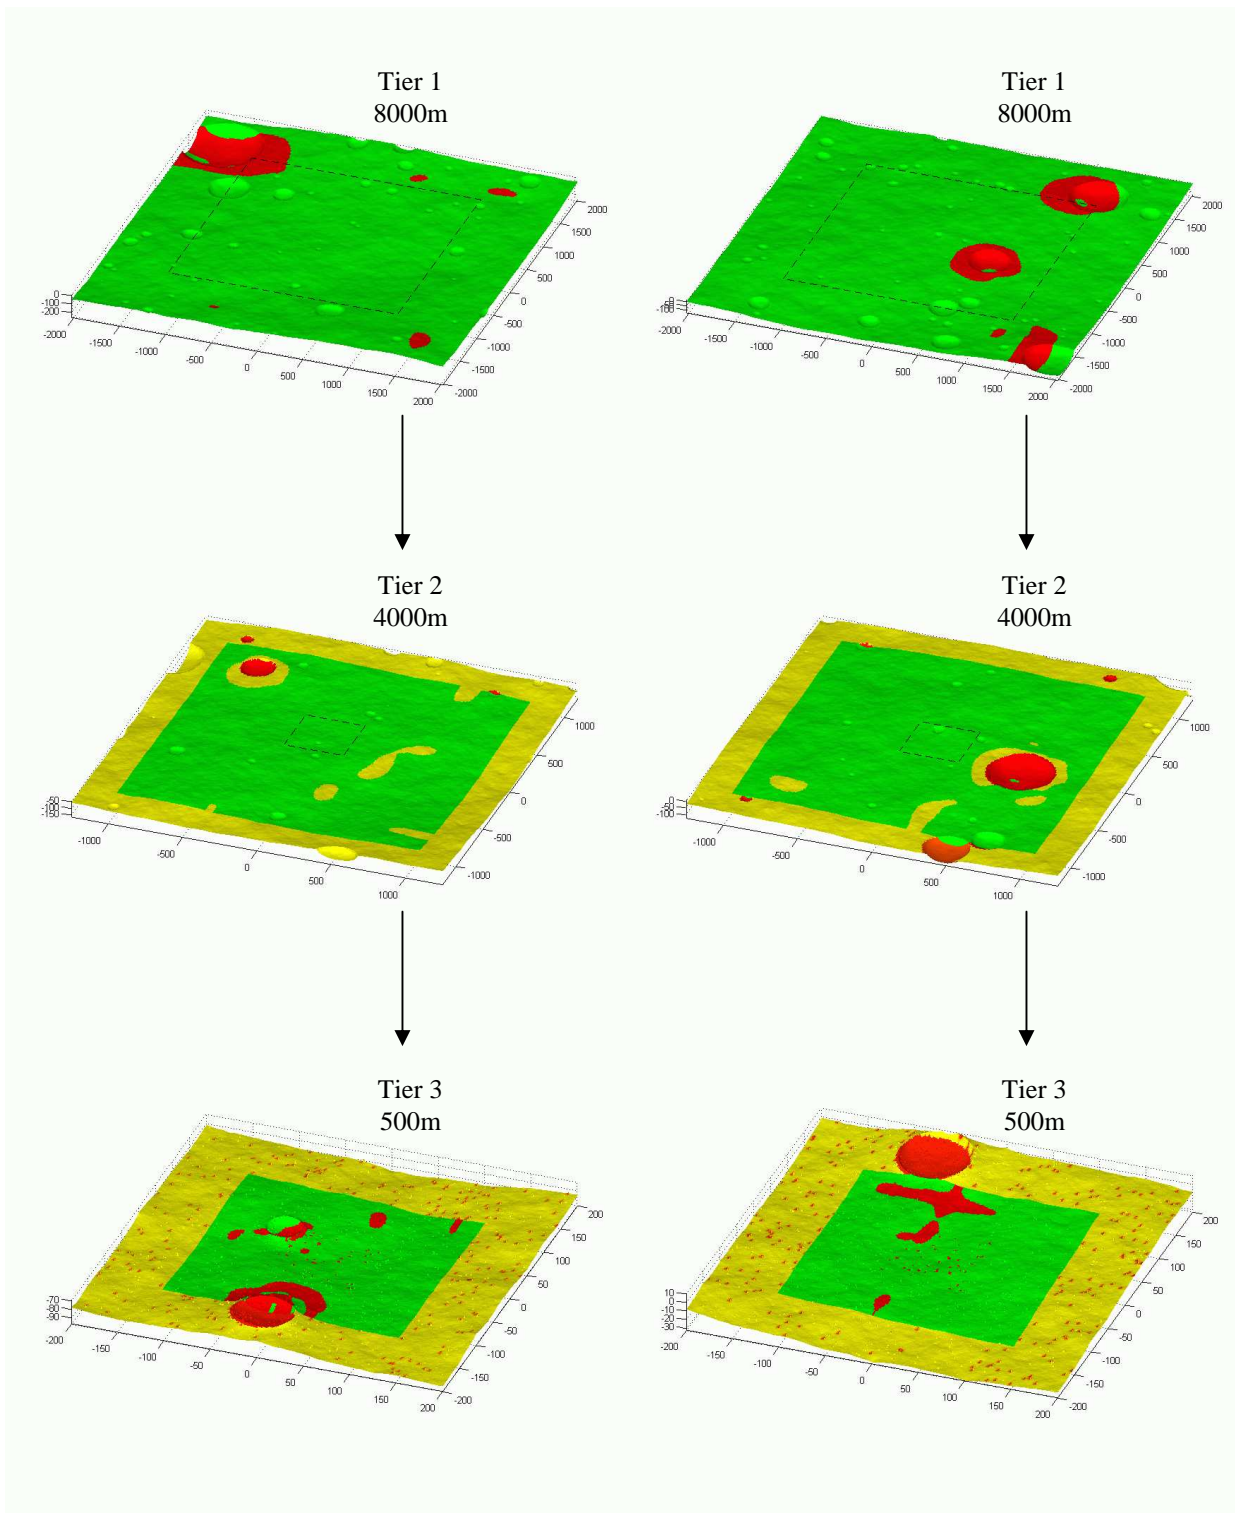
The paper describes a method for determining safe regions of the terrain for an autonomous planetary lander using evidential reasoning. Crater and rock detection algorithms are used to identify unsafe regions of the terrain directly from camera imagery. Classification and Regression Trees are used to determine optimal thresholds that map topographic features extracted from range data to terrain safety. Dempster-Shafer theory is used to derive an evidential terrain safety score that fuses the contributions from each of three sensors (radar, camera, and lidar). The overall approach successfully integrates information from multiple, heterogeneous sensors in a computationally efficient framework with uncertainty management. For future work, plans include a more rigorous comparison between the proposed evidential safety assessment and previously reported Bayesian and fuzzy approaches.

## 8. Acknowledgments

The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology under contract from the National Aeronautics and Space Administration (NASA) with funding from the Mars Technology Program, NASA Science Mission Directorate.

## 9. References

- [1] N. Serrano, M. Bajracharya, A. Howard, and H. Seraji, "A Novel Tiered Sensor Fusion Approach for Terrain Characterization and Safe Landing Assessment," *IEEE Aerospace Conference*, Big Sky, MT, March 2006.
- [2] N. Serrano, "A Bayesian Framework for Landing Site Selection during Autonomous Spacecraft Descent," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, October 2006.
- [3] N. Serrano and H. Seraji, "A Fuzzy Rule-Based Approach to Landing Site Selection," *IEEE International Conference on Robotics and Automation*, Rome, Italy, April 2007.
- [4] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Wadsworth International Group, 1984.
- [5] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [6] Y. Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," *IEEE International Conference on Robotics and Automation*, Barcelona, Spain, April 2005.
- [7] A. Huertas, Y. Cheng, and R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *IEEE Aerospace Conference*, Big Sky, MT, March 2006.
- [8] A. Johnson, A. Klumpp, and A. Wolf, "Lidar-based Hazard Avoidance for Safe Landing on Mars," *AAS/AIAA Space Flight Mechanics Meeting*, Santa Barbara, CA, February 2001.
- [9] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, Second Edition, Wiley-Interscience, 2001.
- [10] J. Balaram, R. Austin, P. Banerjee, T. Bentley, D. Henriquez, B. Martin, E. McMahon, G. Sohl, "DSENDS - A High-Fidelity Dynamics and Spacecraft Simulator for Entry, Descent and Surface Landing," *IEEE Aerospace Conference*, Big Sky, MT, March, 2002.
- [11] A. Fournier, D. Fussell, and L. Carpenter, "Computer Rendering of Stochastic Models," *Communications of the ACM*, 25(6), June 1982, pp.371-384.
- [12] D. E. Bernard and M. P. Golombek, "Crater and Rock Modeling for Mars Landing," *AIAA Space*, Albuquerque, NM, August 2001.



**Figure 5. Evidential safety assessments for two different planetary terrains (left and right columns). The evidential safety assessment is shown for both terrains in tier 1 at 8000m, tier 2 at 4000m, and tier 3 at 500m. The dashed square represents the portion of the terrain visible in the next tier.**