A Novel Tiered Sensor Fusion Approach for Terrain Characterization and Safe Landing Assessment

Navid Serrano, Max Bajracharya, Ayanna Howard*, Homayoun Seraji

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109, USA
firstname.lastname@jpl.nasa.gov

*Human-Automation Systems (HumAnS) Lab
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA
ayanna.howard@ece.gatech.edu

Abstract—This paper presents a novel, tiered sensor fusion methodology for real-time terrain safety assessment. A combination of active and passive sensors, specifically, radar, lidar, and camera, operate in three tiers according to their inherent ranges of operation. Low-level terrain features (e.g. slope, roughness) and high-level terrain features (e.g. hills, craters) are integrated using principles of reasoning under uncertainty. Three methodologies are used to infer landing safety: Fuzzy Reasoning, Probabilistic Reasoning, and Evidential Reasoning. The safe landing predictions from the three fusion engines are consolidated in a subsequent decision fusion stage aimed at combining the strengths of each fusion methodology. Results from simulated spacecraft descents are presented and discussed.[1,2]

TABLE OF CONTENTS

1. INTRODUCTION

NASA uses robotic spacecraft, instrumented with a variety of sensors for science exploration of remote planetary surfaces such as the desolate and rocky terrain of Mars. Spacecraft landers are equipped with necessary sensors and instruments needed to perform a mission at the site where they land. Their utility, though, is dependent on their ability to *safely* land at targeted scientific sites of interest. As such, these robotic spacecraft must land autonomously, without frequent contact with Earth-based mission operators [1].

Typically, engineering criteria established for ensuring mission success are constructed by analyzing terrain characteristics that affect the spacecraft's ability to land safely on a planetary surface within a geo-referenced ellipse of specified uncertainty [2]. Recent images of the surface of Mars have been routinely acquired from orbit by cameras and instruments on the NASA Mars Global Surveyor and Mars Odyssey orbiters. They have acquired images of the landing sites where spacecraft are expected to land. Each spacecraft is navigated to land within a landing ellipse of a size commensurate with the precision of the spacecraft navigation system and with a very high probability of success.

To extract all terrain characteristics that satisfy engineering constraints, a suite of on-board heterogeneous sensors can be used. The data retrieved from multiple sensors must be fused in real-time to ensure safe spacecraft landing. During descent, sensor data are used to analyze the approaching terrain for hazards, and sites within the landing ellipse that are deemed safe for spacecraft landing are selected. This enables any trajectory adjustments to occur in the thruster command sequence in order to minimize the risk to the spacecraft at touchdown.

2. TIERED APPROACH

Three sensors are used to provide measurements of the terrain features during descent: phased array terrain radar, Mars descent imager, and scanning lidar. The characteristics for these sensors are shown in Table 1.

Table 1. Sensor characteristics

	Radar	Camera	Lidar
Min. Range	100m	0m	0.5m
Max. Range	10km	8km	1.5km
Resolution	4.4m @ 1km	0.025m @ 10m	1.25mrad / pixel
FOV	24°	10°	66°
Image size	10×10	1024×1024	100×100

As can be seen from Table 1, the sensors can be naturally grouped in tiers based on their operating ranges. In the first tier, at high altitudes (10km – 8km), only the radar is operational. In the second tier, at medium altitudes (8km –

1km), both the radar and camera are operational. And finally, in the third tier, at low altitudes (1km – touchdown), all three sensors, the radar, camera, and lidar, are operational. Thus, during descent, the terrain is assessed from measurements obtained from one sensor, two sensors, or three sensors, depending on spacecraft altitude, as summarized in Table 2.

Table 2. Tiered sensor grouping

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

Safe landing is the most important factor in landed space exploration missions. If the landing safety assessment is sensor-driven, it is obviously imperative to avoid sensor failure. The use of multiple sensors provides added robustness to safe landing site selection through redundancy. In addition, the sensors chosen for this study provide a complementary set of terrain measurements, which enhance the safe landing assessment.

One of the camera's strong characteristics is its high resolution. However, regions of the terrain with dangerous elevation changes are often undetectable under certain lighting conditions. This limitation is mitigated by the inclusion of the radar, which provides a 2.5-dimensional model of the terrain. Although the radar resolution is coarse, it is sufficient to detect large hazards and changes in elevation; particularly at high altitudes. The inclusion of the lidar, which has a very high resolution, is important at lower altitudes where smaller hazards, such as rocks, must be detected. Finally, the use of both active (radar and lidar) and passive (camera) sensors affords flexibility in dealing with a variety of environmental conditions, which in turn increases the number of feasible landing regions that can be detected.

3. TERRAIN FEATURES

Various terrain features are used to assess landing safety, depending on the operational sensors. The features used by the fusion engines to infer safety must map to known hazards, such as craters, rock fields, ravines, cliffs, steep hills, and so on. One of the unique aspects of the proposed approach is that a combination of low-level terrain features derived directly from the sensor measurements are used together with high-level terrain features which are obtained by landmark detection algorithms. Let $r(x,y)$, $c(x,y)$, $l(x,y)$ represent the raw radar, camera, and lidar sensor measurements at coordinate (x,y) relative to a fixed frame on the terrain surface. The derived terrain features used for landing assessment are summarized below.

Height

This feature is the difference between the radar range $r(x,y)$

and the mean range value $\bar{r}$. The mean range value is computed over all (x,y) and provides an estimate of the terrain surface baseline. The height feature is defined as:

$$f_1(x, y) = \bar{r} - r(x, y) \tag{1}$$

Radar Slope

This feature can be obtained by estimating the slope of a plane locally fitted to the radar range or more simply by calculating the gradient magnitude of the radar range:

$$f_2(x, y) = \sqrt{\left(\frac{\partial r(x, y)}{\partial x}\right)^2 + \left(\frac{\partial r(x, y)}{\partial y}\right)^2} \tag{2}$$

Camera Roughness

A measure of terrain roughness as observed by the camera is obtained by computing the local intensity variance at each pixel:

$$f_3(x, y) = \frac{1}{|W|} \sum_{(i,j)\in W} [c(i, j) - \bar{c}_W]^2 \tag{3}$$

where W is a sub-region of the camera intensity image $c(x,y)$, $|W|$ is the cardinality or number of pixels in W, and $\bar{c}_W$ is the mean intensity in sub-region W. In addition, the feature is smoothed to limit noise effects.

Craters

In order to enhance the terrain characterization, high-level features are also incorporated, that is, features that are derived through additional intelligent processing of the sensor measurements and map to specific landmarks directly associated with landing safety.

The first of such high-level features is the presence of craters. The crater detection algorithm described in [3] is used to determine the shape of craters detected from camera imagery. This is a Boolean feature that represents regions of the terrain where craters are observed. After identifying candidate craters in the camera image based on the sun angle and the cast shadows, the crater detection algorithm fits an ellipse to the crater boundaries. Let $x_{0,i}$, $y_{0,i}$, a_i, b_i, and θ_i be the ellipse center x-coordinate, center y-coordinate, semi-major axis length, semi-minor axis length, and rotation angle, respectively, for the ith detected crater. The Boolean map describing the presence of craters is then

$$f_4(x, y) = \begin{cases} 1, & for\ \mathbf{v}_i^T \mathbf{v}_i \leq 1 \\ 0, & otherwise \end{cases} \tag{4}$$

where

$$\mathbf{v}_i = \begin{bmatrix} \dfrac{x}{a_i} & \dfrac{y}{b_i} \end{bmatrix}^T$$

and x and y are points in a coordinate system rotated by θ_i and translated by $x_{0,i}$ and $y_{0,i}$.

Rocks

The second high-level feature employed is the presence of rocks. Rocks are also detected from camera imagery using the algorithm described in [4]. The algorithm works by first detecting shadows in the image and then determining the size of the candidate rock based on the length of the cast shadow, the known sun angle, and a generic rock model.

Similar to the craters, the presence of rocks is a Boolean feature:

$$f_5(x, y) = \begin{cases} 1, & for\ (x, y) \in R \\ 0, & otherwise \end{cases} \quad (5)$$

where R is the set of pixel locations in the image identified as rocks.

Lidar Slope

The slope of the terrain is estimated by fitting a plane about the local region surrounding every point in the lidar elevation $l(x,y)$. Using the plane model $z = ax + by + c$, the Least Median of Squares (LMedSq) regression method [5] is used to find the parameters of the best fitting plane at each location (x,y). The slope is simply:

$$f_6(x, y) = \sqrt{\hat{a}^2 + \hat{b}^2} \quad (6)$$

where $\hat{a}$ and $\hat{b}$ are the estimates of plane parameters a and b at location (x,y).

Lidar Roughness

Using the plane fitting method described above, the lidar roughness is merely the residual of the lidar elevation $l(x,y)$ and the value of the best fitting plane at (x,y):

$$f_7(x, y) = \left| l(x, y) - (\hat{a}x + \hat{b}y + \hat{c}) \right| \quad (7)$$

where $\hat{a}$, $\hat{b}$, and $\hat{c}$ are the estimated parameters of the best fitting plane at (x,y).

Because the features are derived from different sensors with different operating ranges, only subsets of the full feature set are used during descent. The feature subsets change according to spacecraft altitude, consistent with the tiered approach. A summary of the feature subsets used for terrain characterization is shown in Table 3.

Table 3. Features used in each tier

Tier	Features
1	Height, radar slope
2	Height, radar slope, craters, camera roughness
3	Height, radar slope, rocks, camera roughness, lidar roughness, lidar slope

Although the camera is operational in both tiers 2 and 3, the features derived from the camera are different in these two cases. Craters are only detected in tier 2 and rocks are only detected in tier 3. Examples of the terrain features extracted in each of the three tiers are shown in Figure 1. The features, as shown, have all been registered.

4. SENSOR FUSION

The Entry-Descent-Landing (EDL) operations of a spacecraft occur over a very short period of time, typically on the order of 1-2 minutes. Therefore, the computational speed of any algorithm used for terrain analysis is of utmost importance. The framework employed for selection of safe landing sites combines terrain hazard information extracted from different sensor sources into a global scene description [6]. To obtain hazard information, geological terrain characteristics are identified from sensor data and combined to determine the risk associated with landing within the landing ellipse with the given terrain characteristics.

During spacecraft descent, terrain features are extracted from the operational sensors. These features must then be integrated in order to determine safe landing areas of the terrain. Principles of reasoning under uncertainty are used to assess landing safety based on terrain features observed by the sensors. Three different methods are considered: fuzzy reasoning, probabilistic reasoning, and evidential reasoning.

The same features (discussed in the previous section) are used by each of the three reasoning engines to assess safety. A safety score is determined for each cell in a grid overlaid on the observed terrain. Instead of using a binary, hard-decision safety score (i.e. *safe* vs. *unsafe*), a *multi-level* safety score is used in the continuous-valued interval:

$$s(i, j) \in [0.0, 1.0] \quad (8)$$

where i and j are cell coordinates in the terrain grid. Each reasoning engine will produce a safety score, determined from principles discussed below, in the interval [0.0,1.0]. However, the intermediate values between *highly-unsafe* ($s(i,j)=0.0$) and *highly-safe* ($s(i,j)=1.0$) vary according to the principles of each method as discussed below.
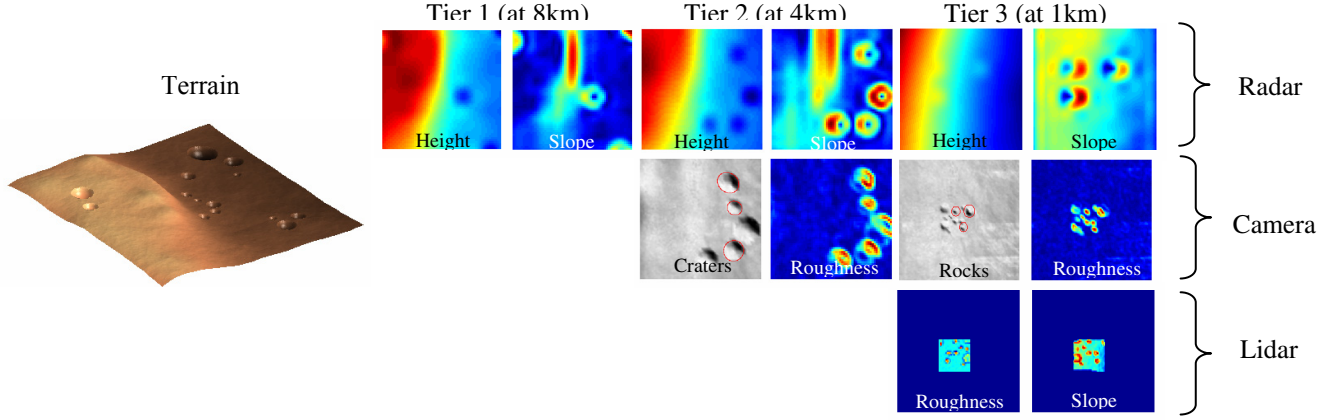
3

Figure 1. Example terrain (left) and extracted features at each tier (right).

Fuzzy Reasoning

To enable real-time implementation, hazard maps are constructed using a computationally efficient reasoning methodology called fuzzy-logic [7]. Fuzzy sets and conditional statements allow the system to manage heuristic rule-based knowledge, imprecise information from sensors, and the uncertainties in the knowledge about the environment. Also, fuzzy rule statements model the human expert's domain knowledge. Fuzzy logic rule evaluation involves only simple arithmetic calculations that can be performed very rapidly. Therefore, the computational time required to create a hazard map is quite manageable for a real-time decision system, making it feasible for landing operations.

At the lowest tier, the data extracted from the heterogeneous sensor suite on-board the spacecraft provide range data that can be converted into an elevation map for extraction of terrain characteristics such as terrain slope and roughness. At the highest tier, the data extracted from the sensor suite include topographic features, such as craters and hills. The hazard map is represented using a grid of cells in which values are represented by fuzzy sets with linguistic labels {HSAFE, MSAFE, MUNSAFE, HUNSAFE}, which stand for *highly-safe*, *moderately-safe*, *moderately-unsafe*, and *highly-unsafe*, respectively. Each cell is associated with a region physically located on the terrain surface. Additional detail on this fuzzy logic methodology is provided in [1].

Individual safety maps are created independently for each contributing sensor on the spacecraft (i.e. radar, camera, and lidar). The safety map for each sensor is obtained by applying the set of rules and membership functions defined for each sensor. These rules are listed in Tables 4, 5, and 6, for the radar, camera, and lidar, respectively. It should be noted that the operators shown in the following tables are fuzzy *t*-norm (AND) and *t*-conorm (OR) operators, not Boolean operators.

Table 4. Radar terrain safety fuzzy rule set

Height	Operator	Slope	Terrain
VLOW	*or*	STEEP	HUNSAFE
LOW	*and*	SLOPED	MUNSAFE
LOW	*and*	FLAT	MSAFE
EVEN	*and*	FLAT	HSAFE
HIGH	*and*	FLAT	MSAFE
HIGH	*and*	SLOPED	MUNSAFE
VHIGH	*or*	STEEP	HUNSAFE

The relation between the linguistic labels for radar slope and terrain safety is quite intuitive. Yet, the inclusion of radar height merits some discussion. By including measurements of terrain elevation, the notion of traversability beyond the landing site is being incorporated. In the proposed approach, landing sites at very high or low elevations are considered unsafe because they typically correspond to hills or craters, which will be avoided. Although local areas within a crater or at the top of a hill may be safe for landing, a rover will not be able to traverse very far beyond the site after landing.

Table 5. Camera terrain safety fuzzy rule set

Hazard	Operator	Roughness	Terrain
PRESENT			HUNSAFE
ABSENT	*and*	VROUGH	HUNSAFE
ABSENT	*and*	ROUGH	MUNSAFE
ABSENT	*and*	SMOOTH	MSAFE
ABSENT	*and*	VSMOOTH	HSAFE

The *Hazard* column in Table 5 refers to either craters or rocks, depending on the tier (see Table 3). If either craters or rocks are detected, the terrain is automatically labeled as *highly-unsafe*. Otherwise, the roughness is used to determine the safety level.

Table 6. Lidar terrain safety fuzzy rule set

Roughness	Operator	Slope	Terrain
ROCKY	*or*	STEEP	HUNSAFE
ROUGH	*and*	SLOPED	MUNSAFE
SMOOTH	*and*	SLOPED	MSAFE
ROUGH	*and*	FLAT	MSAFE
SMOOTH	*and*	FLAT	HSAFE

Tables 4-6 indicate the rules to be used to determine membership in the terrain safety classes. The actual safety score is obtained by defuzzification. The defuzzified safety score derived from each sensor is then fused in such a way that each sensor is allowed to influence the final decision process based on a certainty factor. Fuzzy rules and membership functions are also used to determine the sensor certainty values. The rules are shown in Tables 7 and 8.

Table 7. Radar and lidar certainty fuzzy rule set

Reflectance	Op.	Incidence	Op.	Range	Certainty
DARK					LOW
LIGHT	*and*	SMALL	*and*	NEAR	HIGH
		LARGE			LOW
LIGHT	*and*	SMALL	*and*	DISTANT	HIGH
LIGHT	*and*	SMALL	*and*	FAR	MED

Table 8. Camera certainty fuzzy rule set

Sun Angle	Op.	Angular Motion	Op.	Range	Certainty
		SMALL			LOW
SMALL					LOW
LARGE	*and*	LARGE	*and*	NEAR	MED
LARGE	*and*	LARGE	*and*	DISTANT	HIGH
LARGE	*and*	SMALL	*and*	FAR	HIGH

The final fuzzy-based safety map is thus computed by combining the individual safety maps with sensor certainty values:

$$s_f(i,j) = \frac{\beta_r s_r(i,j) + \beta_c s_c(i,j) + \beta_l s_l(i,j)}{\beta_r + \beta_c + \beta_l} \quad (9)$$

where $s_r(i,j)$, $s_c(i,j)$, $s_l(i,j)$, β_r, β_c, and β_l are the safety scores and certainty values for the radar, camera, and lidar, for the cell (i,j) respectively.

Probabilistic Reasoning

A probabilistic approach to terrain characterization is also considered. Specifically, safety assessment is formulated as a problem of probabilistic inference using Bayesian Networks (BNs) [8]. BNs are directed acyclic graphs (DAG), where the nodes represent variables and the links between nodes represent causal dependence quantified by conditional probabilities. BNs have been used in a variety of fields [8], including data fusion [9]. A BN fully encodes the joint probability of the landing safety and the terrain features:

$$P(S,\mathbf{f}) = P(S)\prod_{i=1}^{N} P(f_i \mid S, f_1,...,f_{i-1}, f_{i+1}...,f_N) \quad (10)$$

where **f** is the full set of N terrain features described earlier. Safety assessment can be formulated as a probabilistic inference problem by application of Bayes' rule:

$$P(S \mid \mathbf{f}) = \frac{P(S,\mathbf{f})}{P(\mathbf{f})} = \frac{P(\mathbf{f} \mid S)P(S)}{P(\mathbf{f})} \quad (11)$$

where S is a discrete random variable representing terrain safety, and **f** is a random vector of terrain features. The joint probability expressed in Eq. (10) assumes the features are all dependent, resulting in an N-dimensional probability distribution. However, computation of the joint probability can be simplified considerably by exploiting conditional independence relationships captured in the graphical structure of the BN. The graphical structure, and in turn the dependence relationships, of a BN can be determined by expert knowledge or from data. In the simplest case, called naïve Bayes [10], the terrain features are all considered conditionally independent given the terrain safety. In this case, the inference problem of Eq. (11) becomes:

$$P(S \mid \mathbf{f}) = \prod_{i=1}^{N} P(f_i \mid S)\frac{P(S)}{P(\mathbf{f})} \quad (12)$$

The features **f** are extracted from terrain location (x,y) and mapped to cell location (i,j) depending on the resolution of the safety grid. A naïve Bayes representation of the terrain features and landing safety are used in this paper. The graphical representation of the BNs is shown in Figure 2.
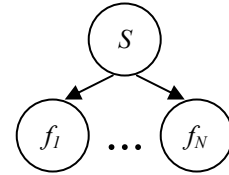

Figure 2. Naïve Bayes graphical representation.

The absence of links between the feature nodes in Figure 2 represents the conditional independence. Such an assumption does not always hold, though for most classification problems it is sufficient [10]. A different naïve Bayes structure is used in each tier, according to the feature subsets shown in Table 3.

Although the terrain safety is modeled as a discrete random variable, S, a continuous-valued safety score can be obtained by using the posterior probability $P(S|\mathbf{f}) \in [0.0,1.0]$, which represents the degree of certainty that the terrain is safe for landing given the observed features. Hence, the probabilistic safety score at cell location (i,j) is simply:

$$s_p(i,j) = P(S \mid \mathbf{f}) \quad (13)$$

Whereas the safety assessment obtained using fuzzy reasoning explicitly incorporates sensor certainty, in the case of probabilistic reasoning, the sensor confidence is

implicit. This confidence is captured by the likelihood $P(S | \mathbf{f})$. This distribution can be calculated analytically with a suitable model or directly from data. Here, a supervised learning approach is used to estimate $P(S | \mathbf{f})$ from data. The features are extracted from a suite of planetary terrains (see section 6 for examples) which have associated safety ground truth. The continuous-valued features are modeled as Gaussian random variables and the parameters of the distributions are computed using Maximum Likelihood (ML) estimation.

Evidential Reasoning

Finally, Dempster-Shafer theory [11] is also considered for combining the sensor measurements. In this approach, terrain safety is derived for each of the sensors and then combined with Dempster's combination rule. Each sensor makes a claim that the terrain is *safe*, *unsafe*, or *unknown* based on its measurements (terrain slope and roughness, or crater and rock detection) and associates a reliability or confidence to this claim. That is, for a claim *C*, each sensor *i* has a reliability *p*. The sensor safety values can then be combined using:

$$m(C) = \frac{\sum_{C_1 \cap C_2 = C} p_1(C_1) p_2(C_2)}{1 - \sum_{C_1 \cap C_2 = \varnothing} p_1(C_1) p_2(C_2)} \qquad (14)$$

where C_i is the claim made by sensor *i*. This provides a degree of belief *m* in the claim. For instance, if two sensors agree that a terrain location is *safe*, the degree of belief in that claim would simply be $1 - (1-p_1)(1-p_2)$. If, on the other hand, they make contradictory claims, the degree of belief in first sensor's claim would be: $p_1(1-p_2) / (1-p_1 p_2)$. The final landing safety is then computed for the fused claims *C*:

$$s_e(i, j) = 1 - \frac{1}{2}[m(C = unsafe) + m(C = safe)] \qquad (15)$$

where *m* is the degree of belief at cell location (*i,j*). In this case, uncertainty is explicitly captured because a sensor's claim that a location is *safe* with some probability does not imply that it is *unsafe* with a complementary degree. Rather, it is a measure of the degree of belief that a location is *safe* and potentially *unknown*, which could mean it is either *safe* or *unsafe*. As a result, a sensor that fails and acts conservatively will not cause the sensed area to be considered *unsafe*, but simply *unknown*, influencing the final combined belief in that area less.

5. DECISION FUSION

Each of the reasoning engines described above arrives at a safety score in the interval [0.0,1.0] that maps to *highly-unsafe* and *highly-safe* at the lower and higher bounds,

respectively. The intermediate values in the range depend on the unique reasoning characteristics of each of the three fusion techniques. Therefore, the distinctive and complementary aspects of each reasoning engine can be exploited and the individual assessments can be combined in a decision fusion step that yields a final safety score.

Min/Max Decision Fusion

The aforementioned fusion engines generally differ in terms of how aggressively or conservatively the safety scores are assessed in various parts of the terrain. Consequently, bounds on the numerical safety score can be obtained by computing the minimum and maximum safety scores for each fusion method at each cell location.

Computing the minimum safety score—equivalent to a logical AND operation—provides a lower bound on the overall safety and can be considered a conservative assessment:

$$s_{min}(i, j) = \min\{s_f(i, j), s_p(i, j), s_e(i, j)\} \qquad (16)$$

Computing the maximum safety score—equivalent to a logical OR operation—provides an upper bound on the overall safety and can be considered an aggressive assessment:

$$s_{max}(i, j) = \max\{s_f(i, j), s_p(i, j), s_e(i, j)\} \qquad (17)$$

Weighted Average Decision Fusion

While the use of min/max operators define the safety bounds, an intermediate safety score can be obtained by linearly combining the predictions of the three fusion methods:

$$s_{avg}(i, j) = w_f s_f(i, j) + w_p s_p(i, j) + w_e s_e(i, j) \qquad (18)$$

where, $w_f + w_p + w_e = 1$. Clearly, the choice of weights w_f, w_p, and w_e is very important. The weights can be set equally ($w_f = w_p = w_e = 1/3$) or can be set to different values in order to bias the result in favor of one or more of the fusion methods.

6. SIMULATION RESULTS

The proposed approach is validated on data obtained from simulated spacecraft descents. DSENDS, a high-fidelity dynamics and spacecraft simulator for entry, descent and landing [12], as well as models of the three sensors (radar, camera, and lidar) are used to extract measurements of the terrain during descent. The DSENDS simulation environment requires topographic models of the terrain and the fusion engines require sufficient terrain features in order to generalize the safety assessment. Thus, Digital Elevation

Maps (DEMs) of various planetary terrains are synthetically generated to simulate spacecraft descent.

The terrain DEMs are created using a fractal model. Hills, craters, and rocks are then added interactively to create a variety of terrains with different landing hazards. The hills are generated using a Gaussian mixture model, the craters are generated using a hemispherical model, and the rocks are generated using randomly perturbed and quantized spheres. Given the suite of terrain DEMs, safety ground truth can be automatically generated based on the location of known hazards. Alternatively, the landing safety can be determined using expert knowledge. A typical synthetic terrain is shown in Figure 1. The terrain consists of an elevated ridge with a steep drop-off towards a valley. Multiple craters of varying sizes are scattered throughout the terrain. Also present, though not visible in the figure, are large boulders of varying sizes.

Using a suite of synthetic terrains, the DSENDS simulation environment is used to produce sensor measurements of the terrain during simulated descents. The sensor measurements are then processed to obtain the terrain features. Given the features, the fusion engines are applied in order to assess safety at various elevations above the surface. Simulation results are shown in Figures 3-6. In each case, the safety assessment is overlaid on the 3D terrain. The safety color-scheme is: red=*highly-unsafe*, orange=*moderately-unsafe*, yellow=*moderately-safe*, and green=*highly-safe*.

The difference between the three reasoning engines is evident from the results shown in Figures 3-5. As can be seen in Figure 3, the fuzzy safety assessment provides more uniformly distributed safety scores as larger portions of the terrain are classified as *moderately-unsafe* and *moderately-safe* compared to the other two methods (Figures 4 and 5). Unfortunately, this leads to one of the craters in the second tier of Figure 3 being classified as *moderately-unsafe*.

The probabilistic safety assessment provides more contrast as most regions of the terrain are classified as *highly-unsafe* or *highly-safe*, as shown in Figure 4. There is less uncertainty in determining the two extremes of *highly-unsafe* and *highly-safe*. As a result, the crater that is missed by the fuzzy reasoning engine in Tier 2 is properly marked as *highly-unsafe* by the probabilistic reasoning. On the other hand, fewer regions are classified in the moderate ranges.

Finally, the results using evidential reasoning can be seen in Figure 5. In this case, the features are mapped to safety levels heuristically for each sensor. Application of Dempster-Shafer's combination rule in Eq. (14) yields a fused safety assessment, as shown in Figure 5. In Figure 5, the progression from Tier 1 through Tier 3 shows that the landing score depends largely on agreement between the sensors. That is, when multiple sensors agree that a particular region of the terrain is either safe or unsafe, the certainty increases. Conversely, when the sensors do not agree, the landing score is set with less certainty.

As can be seen from Figures 3-5, the results obtained from each of the three fusion methods are quite different. This is a desirable outcome in that the differing landing scores can be integrated to produce a more definitive decision. The decision fusion stage allows for the different attributes of each approach to be combined. As discussed earlier, by taking the minimum of the three safety scores, a conservative final assessment is obtained. This result can be used if there is not much *a priori* information about the region where the spacecraft is landing. Computing the maximum safety score provides an aggressive final assessment—one that can be used if *a priori* information suggests the region is mostly safe for landing. By computing a weighted average between the three methods, a more balanced final assessment is obtained. The difference between the three decision fusion approaches is shown in Figure 6.

7. CONCLUSIONS AND FUTURE WORK

A tiered approach to terrain characterization and safe landing assessment is presented. The tiered approach allows for the use of a variety of terrain features derived from different sensors that are operational at various altitudes during descent. Such a framework enhances the safety assessment by providing a richer set of terrain features, as well as providing added robustness. In addition, three different reasoning mechanisms are used to integrate the terrain features and determine landing safety. The reasoning or fusion engines each provide a unique approach to the problem and thus, complementary assessments of landing safety. These complementary assessments are combined in a decision fusion stage. Thus, the final safety score blends the attributes of multiple sensors as well as multiple reasoning engines, which is a unique aspect of the proposed method. Future work will include arbitrary descent paths (which involve more complex data registration techniques) as well as the incorporation of safety predictions over time. Additional means of combining the safety scores in the decision fusion stage will also be investigated.

8. ACKNOWLEDGEMENT

REFERENCES

[1] A. Howard and E. Tunstel, "Using Geospatial Information for Autonomous Systems Control," *Frontiers of Geographic Information Processing*, Eds. Rana and Sharma, Springer Science, November 2005.

[2] M.P. Golombek, R.A. Cook, H.J. Moore, and T.J. Parker, "Selection of the Mars Pathfinder Landing Site," *Journal of Geophysical Research*, 102(E2), pp.3967-3988, February 1997.

[3] Y. Cheng and A. Ansar, "Landmark Based Position Estimation for Pinpoint Landing on Mars," *IEEE International Conference on Robotics and Automation*, April 2005.

[4] A. Huertas, Y. Cheng, and R. Madison, "Passive Imaging Based Multi-cue Hazard Detection for Spacecraft Safe Landing," *This meeting*.

[5] 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*.

[6] A. Howard and H. Seraji, "Multi-Sensor Terrain Classification for Safe Spacecraft Landing," *IEEE Transactions on Aerospace and Electronic Systems*, 40(4), October 2004, pp.1122-1131.

[7] L.A. Zadeh, "Fuzzy Sets", *Information and Control Journal*, vol. 12, pp.338-353, 1965.

[8] J. Pearl, "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference," Morgan Kaufmann, San Mateo, CA, 1988.

[9] L.A. Klein, "Sensor and Data Fusion," SPIE Press, Bellingham, WA, 2004.

[10] N. Friedman, D. Geiger, M. Goldszmidt, "Bayesian Network Classifiers," *Machine Learning*, 29, pp.131-163, 1997.

[11] G. Shafer, "A Mathematical Theory of Evidence," Princeton University Press, 1976.

[12] 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, March, 2002.
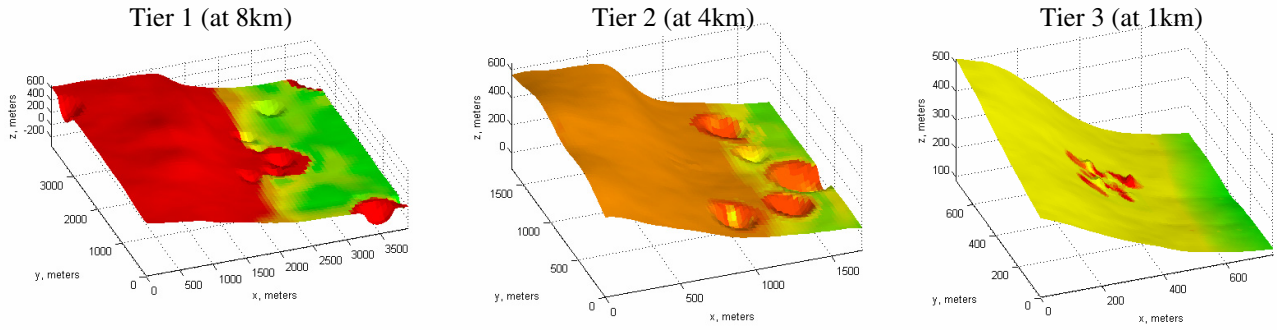
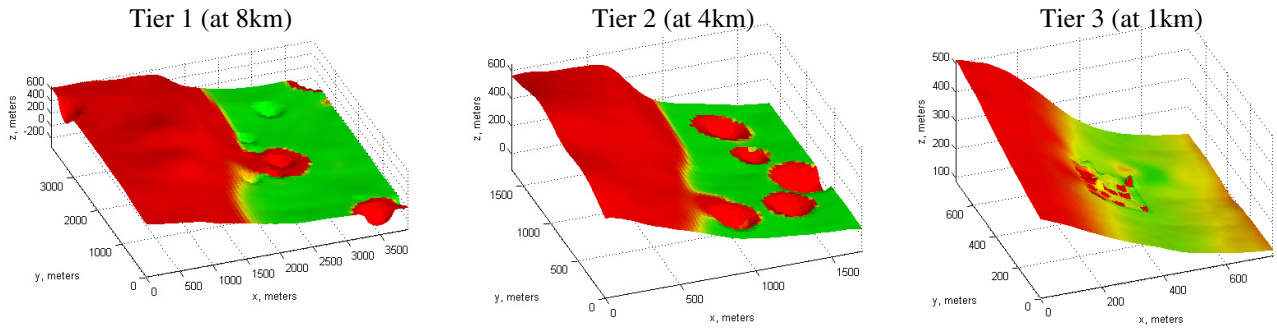**Figure 3**. Terrain safety assessment using fuzzy reasoning.

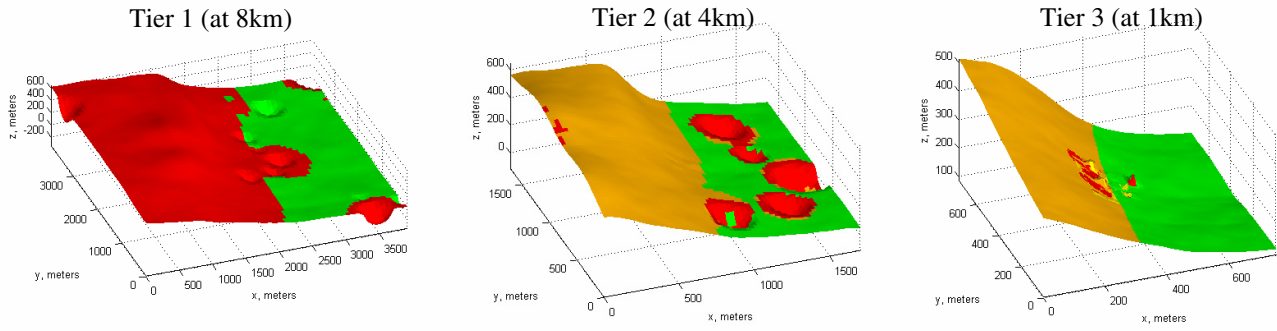**Figure 4**. Terrain safety assessment using probabilistic reasoning.

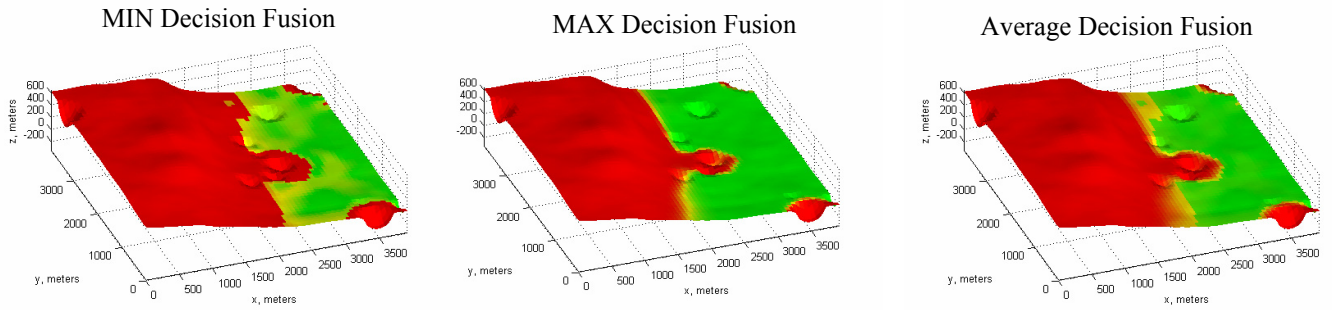**Figure 5**. Terrain safety assessment using evidential reasoning.

Figure 6. Terrain safety assessment in tier 1 (at 8km) after decision fusion.

Navid Serrano is a member of the technical staff in the Computer Vision group at the Jet Propulsion Laboratory, Pasadena, CA. He earned B.S. and M.S. degrees in electrical engineering from the University of California, San Diego, and the Rochester Institute of Technology in 1998 and 2002, respectively. He also holds an M.S. degree in Integrated Media Systems from the University of Southern California. Prior to joining JPL he was a member of the research staff at Eastman Kodak Company, Rochester, NY, where he worked in the areas of image understanding and image enhancement. His current work at JPL is in the area of autonomous safe landing on Mars.

Max Bajracharya is a member of the Mobility and Manipulation Group at the Jet Propulsion Laboratory, Pasadena, CA. His research includes vision-based tracking, manipulation, navigation, and pose estimation for Mars rovers. He is currently the task lead for several tasks focusing on autonomous target approach and instrument placement. Max received his Bachelors and Masters degrees in computer science and electrical engineering from MIT in 2001.

Ayanna Howard is an Associate Professor at the Georgia Institute of Technology. Her area of research is centered around the concept of humanized intelligence, the process of embedding human cognitive capability into the control path of autonomous systems. This work, which addresses issues of autonomous control as well as aspects of interaction with humans and the surrounding environment, has resulted in over 60 written works in a number of projects – from autonomous rover navigation for planetary surface exploration to intelligent terrain assessment algorithms for landing on Mars. To date, her unique accomplishments have been documented in over 12 featured articles - including being named as one of the world's top young innovators of 2003 by the prestigious MIT Technology Review journal and in TIME magazine's "Rise of the Machines" article in 2004.

Homayoun Seraji joined JPL in 1985 as a Senior Member of Technical Staff, following his 13-year academic career, and has been a Supervisor leading and managing a robotics group since 1991. During his tenure at JPL, he has conducted extensive research in: adaptive robot control, control of dexterous robots, contact control, real-time collision avoidance, rule-based robot navigation, and safe spacecraft landing. The outcome of his research in controls and robotics has been published in 93 peer-reviewed journal papers, 112 refereed conference publications, 5 contributed chapters, and has led to 10 patents. In 1996, Dr. Seraji was appointed a Senior Research Scientist at JPL in recognition of his significant individual research contributions in the fields of controls and robotics. He was selected a Fellow of IEEE in 1997 for his contributions to robotic control technology. In 2003, he was recognized as the most-published author in the 20-year history of the Journal of Robotic Systems.