Guidance Augmentation for Reducing Uncertainty in Vision-Based Hazard Mapping during Lunar Landing

Eleanor S. Crane[*,†] **and Stephen M. Rock**[*]
[*]*Department of Aeronautics and Astronautics, Stanford University, Stanford, CA, 94305*
[†]*Moon Express, Inc., NASA Research Park, PO Box 309, Moffett Field, CA 94035*

***Abstract*— A new guidance augmentation scheme, which generates information-seeking trajectory adjustments, is shown to produce improvements in hazard mapping during autonomous lunar landing. This approach utilizes previously developed techniques for detecting hazard objects from images and an Extended Kalman Filter recursive estimation framework in order to create an occupancy grid representation of the hazards in the landing area. The guidance augmentation is driven by a model-predictive scheme which uses predictions of map entropy and fuel usage in order to generate information-seeking acceleration commands which are combined with the targeting capability of Modified Apollo Guidance. An overview of the algorithmic steps required to adjust the trajectory and predict mapping performance and fuel costs are presented. Hazard maps generated using the online information-seeking trajectory adjustments show significant improvement over un-adjusted trajectories where the hazard image data collected is only incidental.**

TABLE OF CONTENTS

1 INTRODUCTION 1
2 BACKGROUND 2
3 ESTIMATION APPROACH 2
4 GUIDANCE APPROACH 3
5 EXPERIMENTAL DESIGN 7
6 RESULTS 8
7 CONCLUSION 11
ACKNOWLEDGMENTS 11
REFERENCES 11
BIOGRAPHY 12

1. INTRODUCTION

Next-generation autonomous planetary landers, which aim to touch down in more varied and variable terrain, will need Hazard Detection and Avoidance (HDA) systems in order to land safely and reliably [1]. By using on-board sensing and mapping techniques, a map of the hazard field at the desired landing area can be created and the final touchdown site selected from it in order to maximize the chance of a successful landing. In challenging terrain, HDA systems can improve the probability of making a successful landing to the level that such missions become feasible. By increasing the landing success rate over diverse terrain, terrain-relative mapping and in-flight landing site selection capability could open up a greater fraction of the Moon's surface for exploration than previously possible.

In a typical HDA scenario, illustrated in Figure 1, the lander descends toward a nominal landing area along some trajectory, computed online. At a pre-specified altitude gate, the landing site may be redesignated based on terrain information, often in the form of a hazard cost map, available at that point; the lander then executes the actual avoidance maneuver in preparation for touchdown at the new site. Previous development by the authors introduced a recursive estimation framework for hazardous rocks using images from a monocular camera, the final output of which is a probabilistic map of the hazards in the area which explicitly and rigorously accounts its for estimation uncertainty [2]. This hazard cost map can then be used at the decision point in order to make the final site selection.

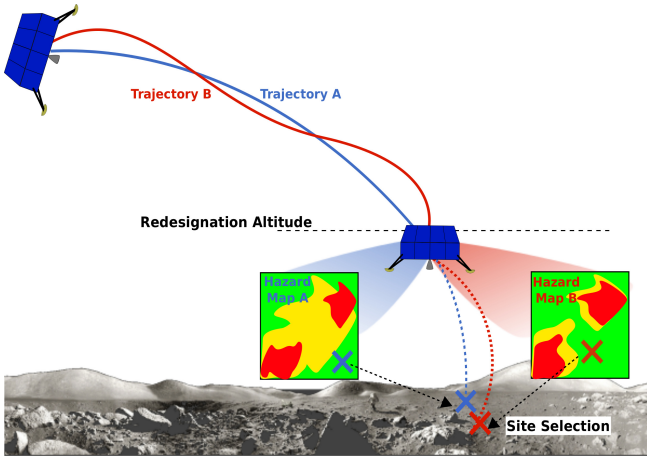

Figure 1. HDA scenario: hazard maps are path dependent

This work seeks to actively improve the hazard maps created during descent by making small modifications to the landing trajectory in order to increase the quality and quantity of information collected about the targeted terrain. Previous work by the authors demonstrated that the trajectory flown during descent can substantially lower the level of hazard map uncertainty at the site selection cutoff altitude, but used only heuristic profiles for comparison [2]. This paper presents a systematic methodology for generating such trajectories during the descent.

A new information-seeking guidance augmentation scheme, which makes small modifications to the descent trajectory, is here shown to produce reduced uncertainty in hazard maps created during autonomous lunar landing. The guidance scheme seeks to decrease hazard map uncertainty by minimizing occupancy grid entropy, used as a measure of uncertainty, at a specified cutoff altitude. This reduction is accomplished by adjusting the descent trajectory in order to improve imaging, at the cost of increased fuel usage.

This paper presents a model predictive controller (MPC)

978-1-4673-1813-6/13/$31.00 ©2013 IEEE.
[1] IEEEAC Paper #2272, Version 2, Updated 01/26/2013.

which generates the guidance augmentation acceleration commands. The derivation of the controller and the information theoretic models which drive the predictions are detailed, as well as the cost function which the controller seeks to minimize. Hazard maps generated in simulation using the MPC augmentation scheme are then presented in order to show that map uncertainty is appreciably reduced. Trajectories flown and the characteristics of information-rich trajectories are detailed, as well as the fuel costs associated with the information gain achieved. In addition to the reduction in map uncertainty, an increase in the quantity of terrain correctly identified as landable is presented. In one case study, for a 14% increase in fuel cost, the MPC guidance augmentation produces a map with an additional 33% of the terrain identified as landable, with a total success rate of 97% at correct identification. These results show that the MPC scheme produces significant improvement over maps generated through un-adjusted trajectories where the hazard image data collected is only incidental.

2. BACKGROUND

The sensor chosen for this task is motivated by the challenges of small spacecraft, which have not only a low tolerance to rocks and other hazards, but also limited power and mass resources with which to perform HDA. To this end, a monocular, vision-based system is proposed for such landers. A passive imaging system has the advantage of being a low power solution, but its utility is limited to the lunar day. However, vision-based techniques can be used very effectively on the lunar surface, where it is illuminated, due to high-contrast lighting and the lack of atmospheric distortion [3]. Active range-sensing methods for HDA, such as LIDAR, have also been extensively studied. Several LIDAR-based systems have been proposed and developed for use on the Moon [4] and Mars [5] to varying levels of technology readiness, and the link between measurement quality and pose has also been noted for these systems [6]. While LIDAR is well-suited to the task of hazard detection, the high power required to operate such a system make it, at present, challenging to house on a small lander with limited resources.

Pursuant to the monocular camera HDA task, shadow-based vision techniques [7], [8] for hazard detection have been developed and tested by the robotics division at the Jet Propulsion Laboratory (JPL). This work has been tested on orbital, aerial, and laboratory imagery, with site selection demonstrated. The published work has focused on the image processing necessary to do the rock detection, but not on an estimation framework to leverage these measurements throughout the descent. Previous work by the authors has built on the hazard detection algorithms developed and proved at JPL, and incorporated them into a larger estimation and guidance architecture for use in the information-seeking system presented here. The framework presented in this paper has been tailored for use with this shadow detection technique, but the larger idea of leveraging the trajectory design to improve a vision-based hazard detection system is not dependent on this particular image processing technique and could be adapted to a different detection algorithm.

The estimation part of the HDA task involves collecting and organizing hazard data, then selecting a new site relative to those hazards. Methods for hazard data collection are largely dependent on the chosen sensing method. Once hazards have been detected, they are most often organized into a map representation and extended to include other costs beyond survival risk. Vision-based methodologies for all types of craft often perform landing site risk evaluations in the frame of the image collected [9], which implies that vision is the only sensor and that only one image is used to make the determination. Multi-sensor fusion approaches have been developed to combine the information from multiple sensors, depending on their heuristic strengths, and evaluated using fuzzy logic criteria to assess risk and suitability [10]. This technique has been extended to become a probabilistic inference network [11], but the models used to represent the inference are simple Gaussians, and can be further extended. Work in the areas of purely active sensing have tended towards constructing elevation maps before extracting risk metrics from the sensor data [5], [12]. Regardless of the hazard sensing method, making the relevant trades to select a final landing site have been studied at length [13] and can be decoupled from the choice of sensor, so long as the hazard map captures the terrain risk associate with a particular set of sites.

Within the robotics community, work with gaze control, or next-best-view problems, for monocular vision systems has been addressed using, among other techniques, information-seeking algorithms which attempt to maximize mutual information using measurement models. In particular, in humanoid robotics, which attempt to mimic human perception, occupancy grids and mutual information are frequently used. The sensor model is often a simple Gaussian fit to the expected camera field of view [14], [15], and is used to decrement occupancy grid entropy and evaluate various control actions. While the link between map uncertainty and robot pose is noted [15], these works usually consider the issue of gaze control separately and after that of path planning: the algorithms seek to maximize information using only the camera pointing and not the robot position. Given the decoupled nature of this approach, a one-step look ahead method is most often used, as there is no down side to a greedy search. However, given that the challenge of HDA for landing craft involves constructing a hazard map using resolution-limited measurements, a greedy approach could prove insufficient.

3. ESTIMATION APPROACH

The information-seeking guidance approach requires several algorithms to produce a map of the hazard field and quantify the uncertainty: first the terrain information must be extracted from images. This information is then put through a recursive estimation system, from which an occupancy grid map of the terrain is created and evolved for use in landing site selection.

Image Processing

Terrain information is collected using images from a monocular camera and processed using a rock detection algorithm developed at JPL, which infers the size of the object based on the size of its shadow [7]. The images are adjusted to increase contrast and segmented into regions of shadows using entropy maximization. Shadow contours are traced and, using estimates of the sun azimuth direction, the length of these shadows are measured. A hemispherical rock model is used to locate the highest point, and center, of the rock as the point on the shadow contour closest to the sun direction. By using the known sun elevation angle from the ground plane, the rock height can be estimated. More complicated rock models have been developed [7], but the hemispherical model is taken as a simple but conservative model based on the assumption that most rocks are wider than they are tall.

The rock detection process is illustrated in Figure 2 using a test image taken in a laboratory setup. The original image is shown on the left, with the shadow outlines, extracted using a Canny edge detector on the gamma-corrected image. The individual rock locations and radii, in pixels, inferred from these shadow outlines are shown on the image on the right. Those measured rocks which correspond to actual rocks are shown in green. Those measurements shown in red have been labeled as false positives. However, while the red crosses may be false positive sightings of rocks, they are generally located in shadowed regions of the surface. The ridges which produced these shadows are generally not safe landing areas, and so are not considered an issue with the method. Not all radii were measured accurately from this single image, but successive images from multiple view-points enables more accurate estimate when combined in a filter.

Figure 2. Intermediate steps in rock detection image processing

Estimation: EKF/Occupancy Grids

A system for recursively estimating the spacecraft and hazard states, in the form of a sequential, Multiplicative Extended Kalman Filter (EKF) [16], has been developed. The spacecraft motion is propagated using an inertial measurement unit (IMU) and is updated using measurements from an altimeter and the monocular camera ($1hz$ image updates). The filter estimates the spacecraft translational states and the location and radii of rock hazards in a local East-North-Up (ENUP) frame. The attitude state is tracked via a three element Gibbs vector and a reference quaternion. Accelerometer and gyro biases are tracked inside the filter for use in spacecraft motion updates. The internal state vector is given by

$$\vec{x} = \begin{bmatrix} \vec{r}_{\text{sc}} & \vec{v}_{\text{sc}} & \vec{b}_{\text{accel}} & \vec{a}_{\text{g}} & \vec{b}_{\text{gyro}} & x_j & y_j & r_j \end{bmatrix} \quad (1)$$

where $\vec{r}_{\text{sc}}$ and $\vec{v}_{\text{sc}}$ are the spacecraft position and velocity in the ENUP frame, $\vec{b}$ are the IMU biases, and $\vec{a}_{\text{g}}$ is the Gibbs vector attitude in the body frame. The reference quaternion is given by $\vec{q}_{\text{ref}}$ and updated with the filter. The hazard objects are appended to the end of the state as they are discovered, with the East, x_j, and North, y_j, positions and estimated radius, r_j for each hazard j. Each hazard is modeled as a hemisphere, regardless of the true shape of the rock imaged, and thus the estimated shape is described by only one parameter.

In order to perform landing site evaluations, the estimated hazard state is transformed from discrete objects into a dense map of the terrain using an occupancy grid representation [17]. This representation shows the probability of a particular map location containing a hazardous object, $p(x_i, y_i)$, where the value 0 corresponds to the grid point i being clear and 1 similarly corresponds the site being occupied by a hazard, and thus unsuitable for landing.

The probability of occupancy is based on both the current best estimate of the hazard size and location, but also the uncertainty about that estimate, in the form of the Gaussian covariance which is output by the Kalman filter framework. The probabilistic distribution of rock sizes is also taken into account using the Golombek model [18], which describes the fractional area of terrain covered by rocks of a certain size or larger, parameterized by the total coverage fraction. In this way, measurements of rocks detected by the image processing techniques can be included to increase the probability of occupancy where appropriate, and clear image measurements can be utilized to decrease the probability of occupancy where no rocks have been detected. In both cases, the probability of incorrect detection is included based on image processing success rates. Further details about the creation of such an occupancy grid map is available in previous work by the authors [2].

Example occupancy grid representations of the same rock field are shown in Figure 3. The contours of probability are shown in the map on the left, with the rock field overlaid in black. This map shows good convergence to the true rock field as areas of high probability are clustered around the rocks. The occupancy grid still shows some areas of uncertainty around each rock, which can be measured in the size of the regions of middling probability which surround them. Blue areas have been identified as being clear of hazards. The map on the right shows the same data with thresholds applied to the continuous probabilities to label each grid point as either 'clear' or 'occupied' based on the probability of occupancy.

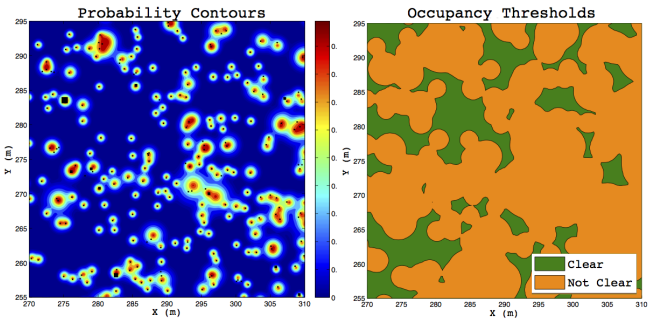

Figure 3. Occupancy grid: contours and thresholds

4. GUIDANCE APPROACH

The information-seeking guidance system works in tandem with the estimation system: the guidance chooses a trajectory which trades predicted hazard map uncertainty against fuel usage. As the lander collects images of the target area, these predictions are improved and further control action can provide increasingly detailed knowledge of the terrain. While this information is collected, the lander must still be making progress towards the target area. These two functionalities are implemented using a baseline guidance algorithm and the information-seeking augmentation algorithm.

The baseline guidance scheme, which delivers the lander to the target area using a nominal acceleration profile, is augmented with small acceleration commands which change the lander pose (position and attitude) at downrange imaging opportunities. These acceleration commands are generated using a discretized search space around the nominal profile, and evaluated using a cost function which trades decreases in hazard map uncertainty against the fuel consumed to move the trajectory away from the nominal.

The model-predictive controller, which drives guidance augmentation, uses models to relate spacecraft pose at imaging opportunities to expected information gain. The relationships between pose and image information are examined in this section, formalized into explicit predictive models, and then harnessed by the MPC framework and traded against fuel using a cost function.

Baseline Guidance: Modified Apollo

The guidance augmentation system described here has been developed and used in tandem with the Modified Apollo Guidance (MAG) approach [19]. This scheme constructs a quadratic acceleration profile through the solution of a two-point boundary value problem which takes the lander from the current position to the desired landing state within a specified time. The calculation of the time-to-go, t_{go}, is set in this framework by limiting the spacecraft jerk during the descent.

The guidance calculation is run online, at a low rate, during the descent: accumulated navigation and pointing errors can then be flown out using these guidance updates. The current spacecraft (sc) and touchdown states (TD) are given by:

$$\vec{x}_{\text{init}} = [\vec{r}_{\text{sc}} \quad \vec{v}_{\text{sc}}] \tag{2}$$

$$\vec{x}_{\text{final}} = [\vec{r}_{\text{TD}} \quad \vec{v}_{\text{TD}} \quad \vec{a}_{\text{TD}} \quad t_{\text{go}}] \tag{3}$$

The solution of the boundary value problem given by these conditions yields a set of polynomial coefficients, $\vec{C}$. The acceleration profile which takes the spacecraft from the current position to the touchdown site, calculated for each axis, is given by

$$\vec{a}(t)_{\text{comm}} = C_0 + C_1 t + C_2 t^2 \tag{4}$$

which is valid on the interval $t \in [0 \quad t_{\text{go}}]$. One useful property of this method is that at every computation of the guidance solution, the spacecraft has a predictable acceleration profile which is will follow until the next computation of the guidance.

Information-Seeking Augmentation

The information-seeking augmentation system calculates acceleration adjustments which modify the trajectory with the goal of decreasing map uncertainty. The goal of the guidance augmentation approach is to harness the targeting capability of the traditional guidance scheme and provide only small adjustments to the trajectory while still making progress towards the goal site. The results presented in this paper are based the Modified Apollo Guidance described above, but it could be possible to augment other guidance schemes as well.

Image Pose Parameters which Affect Information

The basis for calculating the information-seeking inputs must be a strong understanding of the relationship between lander pose and the quality of information collected from an image taken at that pose. These properties should hold regardless of what the actual rock formations are in the target area; the goal of this section is to establish what lander pose parameters produce images with high information potential.

Field of View—This work assumes a body-fixed, nadir pointing camera, which means that the area of the ground imaged depends entirely on the spacecraft position and pointing direction. In order to collect hazard information about a site, obviously, that site must be visible within the camera field of view and so the spacecraft must be pointed towards the target area. A pan-tilt camera would increase the flexibility of the system, but would not completely decouple the link between spacecraft attitude and camera field of view due to tilt limitations and possible obstructions.

Resolution—For a fixed focal length camera, the resolution of the image captured is dependent on the altitude at which the image is taken. By increasing the resolution of the image, the size of the rocks as measured via the length of their shadows becomes more accurate and reduces the probability that a small rock is missed due to noise rejection. It makes intuitive sense that to take a better picture, the lander should be closer to the surface.

Camera View Angle—The relative location of the lander with respect to the sun direction and the area of interest also affects the information content, as the image processing technique relies on detecting the shadow cast by an object to determine its size. Crucially, if the lander is on the opposite side of a rock from its shadow, the shadow can be occluded by the rock itself; any image measurements taken of the rock size from that vantage point could therefore be biased.

In order to examine the effects of lateral position on the image measurement, an angular representation of the relevant geometry is introduced such the parameter can be examined independent of the lander altitude. The angle γ is defined between the direction of illumination and the vector from the lander to the target rock, as shown in Figure 14. For γ less than 90°, the lander camera is viewing the rock from its shadowed side. For γ greater than 90°, the lander camera is viewing the rock from its illuminated side, which produces occluded views.

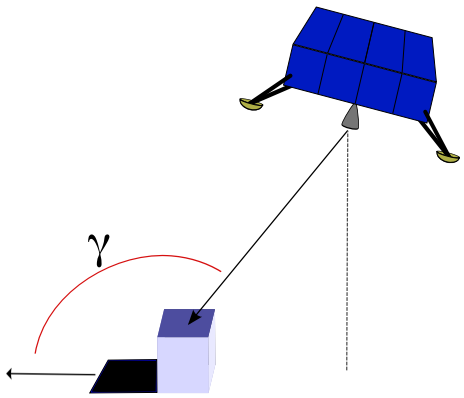

Figure 4. Information sources: view angle

The effect of gamma angle is demonstrated in the following experiment shown in Figure 5: a single object (rectangular prism) is imaged at two different γ angles: 81° and 100°. The shadow lengths are measured from the resulting images. It can be seen from the image on the right that the shadow is occluded, and the visible length which can be measured is shorter than the total length captured in the left image. The length of the shadow is underestimated by 13 pixels, and thus, in this example case, the height of the rock is underestimated by 13%, not including any pixel quantization errors.

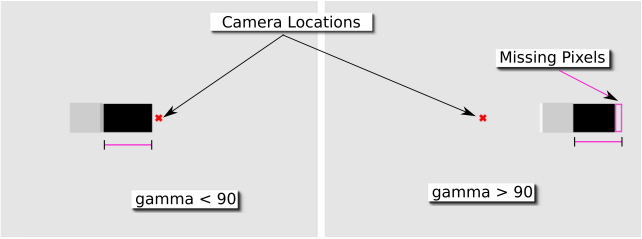

Figure 5. (a - left) Image taken at $\gamma = 81°$, (b - right) Image taken at $\gamma = 100°$

The above experiment was conducted at low altitude, but the effect continues at higher altitudes. Figure 6 shows a plot of rock height measurement error from images of the same $1m$ object, taken at various altitudes and γ angles. The underestimation of the object size continues at higher altitudes, plateauing around 15% height error. At low altitudes, perspective distortion of the object can also increase the measurement error, which accounts for the increase in error evident at very low altitudes.

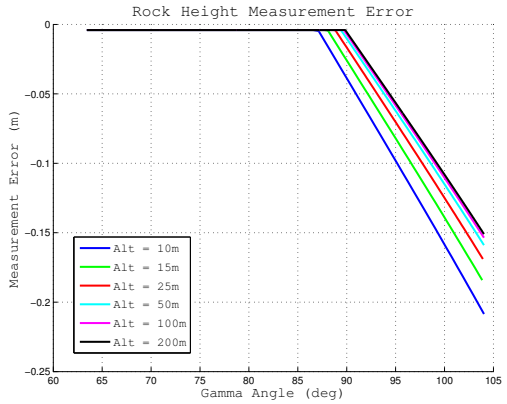

Figure 6. Rock height measurement error as a function of γ

The above experiment was conducted using a rectangular prism as the imaging target, but natural rocks are of such shapes that the progression of error with γ is not nearly as predictable, and therefore cannot easily be compensated for. Rather than adjusting the measured value of shadow length when γ is greater than $90°$, the uncertainty associated with that measurement is increased in the EKF. However, the best measurements of shadow length and rock height are available from the shadowed side of the rock.

Map Entropy/Image Information Prediction

Using the relationships defined above, the guidance augmentation scheme can predict which candidate trajectories will be able to collect images that have a high hazard information potential and then evaluate them based on the expected reduction in map uncertainty. Explicit models for image potential information and map uncertainty reduction are constructed in order to assign a value to each candidate trajectory.

Map uncertainty is measured via the information-theoretic measure of entropy, as applied to the hazard occupancy grid map. The projected vehicle poses are obtained by integrating the candidate trajectory acceleration profile; since the actual image is not predictable, particularly before the landing zone is reached, the links between pose and information quality explored above must be captured explicitly. Therefore, the goal of this section is to construct a model which links spacecraft pose with the change in entropy of the landing zone occupancy grid map.

Hazard map entropy prediction begins with the occupancy grid map constructed up to the current time, given by $P(t_k)$. The current entropy of the map, H_k, is given as the sum of the entropy of individual grid cells, i, for each binary random variable [20]:

$$\begin{aligned} H_k &= \sum_i H\Big(p(x_i, y_i)\Big) \\ &= \sum_i -\Big[p_i \log\left(p_i\right) + (1-p_i)\log\left((1-p_i)\right)\Big] \end{aligned} \tag{5}$$

At the next imaging opportunity, the lander altitude and attitude, along with an internal camera model, are combined to determine what area is within the camera field of view. Given the limited size of the occupancy grid, each ground point is put through the projective model at the current spacecraft pose to determine if it is within the current image boundaries. If the grid point will be imaged, the predicted probability is updated. The image update, since the actual image content cannot be predicted, is modeled to provide confirmation of the current occupancy state: the probability of occupancy is pushed closer to either 0 or 1. During the prediction, no new hazard objects can be added to the map, which limits the effective resolution of the predicted map. The estimation system must provide updated maps to the guidance in order to converge to the true hazard field.

The convergence is modeled mathematically using an additive probability, which changes sign based on the current occupancy. If the probability of a grid point is higher than some *occupancy_threshold*, it is assumed that this grid point actually contains a hazardous object. If the grid point probability is below the *clear_threshold* value, it is assumed to be clear at the current resolution. The additive change in probability due to an expected image is given by the value δp. The algorithm for predicting the evolution of the occupancy grid is given in Algorithm 1.

The *occupancy_threshold* value is biased high, as it is assumed that when objects are added to the map their size is overestimated due to high covariance, and so areas of middling probability are likely to be reduced as the filter converges on the true rock field. The *clear_threshold* value is generally set low, and is based on the prior probability of occupancy. Below this threshold, the probability is set using the Golombek model [18], which is given by the cumulative distribution function, F_G, and parameterized by the total rock coverage fraction, k_r, and the diameter of rocks which would harm the lander, D_{min}, and the diameter currently resolvable by the camera, D_{cam}.

The parameter δp_i is then set based on the lander view angle γ with respect to the particular grid point i. If the γ angle is less than $90°$, the measurement is modeled to provide a larger reduction in the hazard occupancy grid; measurements with the γ angle greater than $90°$ are treated as less helpful in the goal of reducing hazard map uncertainty based on the possibility of shadow occlusion. The numerical values for δp have been chosen based on the expected rate of convergence

Algorithm 1 Occupancy Grid Prediction

procedure PREDICT_OCC($P_k, image_bounds(k+1)$)
  for all points i in occupancy grid **do**
    if point $i \in$ image_bounds$(k+1)$ **then**
      Calculate $\gamma_i(\hat{r}_{sc}(k+1), x_i, y_i)$
      if $\gamma_i < 90^\circ$ **then**
        $\delta p = 0.05$
      else
        $\delta p = 0.05/3$
      end if
      if $P_k(i) >$ occupancy_threshold **then**
        $P_{k+1}(i) = P_k(i) + \delta p$
      else if $P_k(i) >$ clear_threshold **then**
        $P_{k+1}(i) = P_k(i) - \delta p$
      else
        $P_{k+1}(i) = F_G(D_{\text{cam}}, k_{\text{r}}) - F_G(D_{\text{min}}, k_{\text{r}})$
      end if
    else
      $P_{k+1}(i) = P_k(i)$
    end if
  end for
  return P_{k+1}
end procedure

of the hazard objects in the EKF estimation system.

The probability of occupancy is updated at all visible grid points at each predicted imaging opportunity. The total map entropy is calculated using the predicted probability map at the cutoff altitude for use by the guidance augmentation system in evaluating the candidate trajectories.

Model-Predictive Approach

The model-predictive approach uses the models developed above for entropy prediction in order to drive down map uncertainty at a specified cutoff altitude. This approach differs from standard information-seeking control approaches commonly used in robotic mapping, as the constraints of the landing application call for a longer time horizon for the prediction. The use of a cost function to trade information against fuel expenditure is common to both approaches. The standard approach is outlined, and the model-predictive approach is detailed; the cost function used for trajectory and control action selection is also developed.

The goal of the standard information-seeking guidance approach is to maximize the value of mutual information, that is, the information gained about the environment from the next pose [20]. Since the actual value of the measurement, z, cannot be predicted, the expected value of the entropy at the next pose is used instead in order to evaluate the potential control choice, u, to be taken at state x to get to x' and take measurement z.

$$H(x'|u) \approx E_z\big[H(x'|u,z)\big] \tag{6}$$

Using the value of map entropy at the current state, $H(x)$, the mutual information, I, gained by taking control action u is then given by

$$I(u) = H(x) - E_z\big[H(x'|u,z)\big] \tag{7}$$

Potential control actions are then evaluated using a cost function which weighs the expected information gain against a cost associated with taking such a control action, r, which is defined to be negative. In many robotic mapping applications, a one-step lookahead algorithm is used and a greedy policy, π_g, is implemented:

$$\pi_g = \arg\max_u \left(\alpha I(u) + \int r(x,u)dx\right) \tag{8}$$

The model predictive approach proposed for use in lunar landing applications extends this idea to use a long time-horizon prediction for map entropy and fuel usage, while only the first control action is actually taken before another computation of the guidance algorithm. A discretized search space of acceleration adjustment commands is constructed at each guidance calculation in order to limit the computational burden on the flight system. The MPC scheme is summarized in the following cartoon, shown in Figure 7.

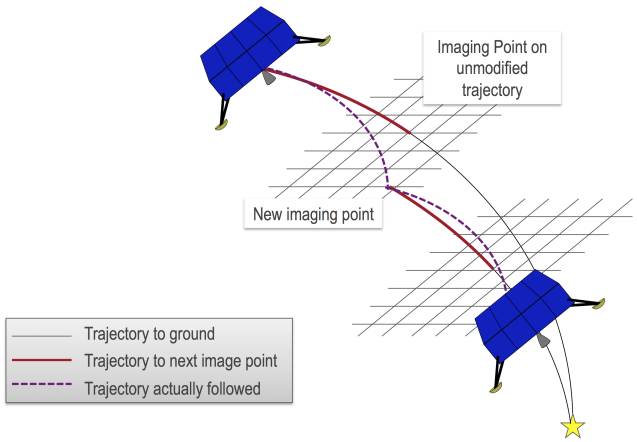

Figure 7. Model-predictive scheme with discretized search space

At each computation of the guidance, the Modified Apollo coefficients which create the acceleration profile are calculated to guide the spacecraft to the goal location. The spacecraft pose is then calculated at the next imaging opportunity by propagating the current state forward in time using the Apollo profile. This position is then allowed to vary on a grid, up to a specified amount away from the nominal position. The magnitude of the position change is decreased as the lander approaches the surface in order to limit large, costly deviations from the nominal trajectory at low altitudes. The additional acceleration, taken to be constant over the required time, is then calculated to create a set of potential guidance augmentation commands, $\delta\vec{a}$.

Similar to the standard information-seeking algorithms described above, each candidate control choice is evaluated based on the information and control cost, in this case propellant usage, to determine which is the optimal choice. For each possible control action, the hazard map entropy at the next imaging opportunity and the fuel used to achieve this pose is predicted, assuming that this action was taken. This modified pose shall be called $\vec{x}'$.

From the position and velocity state at $\vec{x}'$, new MAG polynomials, $\vec{C}'$, are calculated which would take the lander from this new pose to the goal location. The spacecraft state is

propagated forward in time from the pose $\vec{x}'$ using $\vec{C}'$ only to generate the acceleration profile; no further adjustments are made to this candidate trajectory. At each imaging opportunity, the map entropy prediction is updated to include the expected image measurement. The propellant required to follow this candidate trajectory is predicted as well.

The candidate trajectory is propagated forward in time until the desired map cutoff altitude, alt_f, is reached. The predicted map entropy at this altitude, $\hat{H}_{map}(\text{alt}_\text{f})$ and the fuel cost which would be incurred, $\hat{m}_\text{fuel}$, are combined into a cost function which is used to evaluate each candidate trajectory.

The cost function consists of these two terms and an image regularization term, which rewards imaging the target area. This regularization is calculated as the fraction of the landing area which is imaged in a given opportunity. In this way, if the predicted entropy is the equal for two candidate trajectories, preference is given to the one which provides more images of the site. This regularization term is calculated over all imaging opportunities using the points in the occupancy grid, i:

$$\hat{I}_\text{traj} = \Big(\sum_\text{images} \sum_i v(x_i, y_i) \Big) \Big/ \Big(\sum_\text{images} \sum_i 1 \Big) \tag{9}$$

$$v(x_i, y_i) = \begin{cases} 1 & \text{when point } i \text{ in image,} \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

The final cost function is given as

$$J = \alpha \hat{H}_\text{map}(\text{alt}_\text{f}) - \beta \hat{I}_\text{traj} + \hat{m}_\text{fuel} \tag{11}$$

Given the limited size of the search space, all possible control actions within that space are evaluated and the minimum cost value is chosen for implementation. The values α and β are weighting factors which can tune the tradeoff between entropy, imaging, and fuel usage, and are set using prior knowledge of the total rock coverage at the target landing zone.

This cost function shows the same form as the standard information-seeking algorithm outlined above in Equation 8, but the evaluation point is N steps away, rather than one step. The evaluation of the cost at the controller cutoff altitude, chosen as $150m$ in this study, is used to select between all the candidate pose adjustments $\vec{x}'$. The final acceleration command to the vehicle is calculated using the command generated using the original Modified Apollo algorithm, as given in Equation 4, plus the adjustment value $\delta\vec{(a)}$ required to achieve the pose $\vec{x}'$.

5. Experimental Design

Sim Environment Overview

In order to demonstrate the performance of the guidance augmentation system, lander descents onto lunar terrain models are simulated, and the resulting maps compared. This experiment was implemented in a high-fidelity lunar landing simulation environment which has been developed in MATLAB™ with the capability of generating and processing synthetic images of terrain in order to incorporate image measurements into a dynamic, closed-loop simulation of a landing scenario. The simulation environment employs a six degree-of-freedom dynamics model, driven by lunar environment and actuator models. In addition to the image measurements, an IMU and laser altimeter are modeled. All these sensor measurements are incorporated in an implementation of the online estimation system described in Section 3.

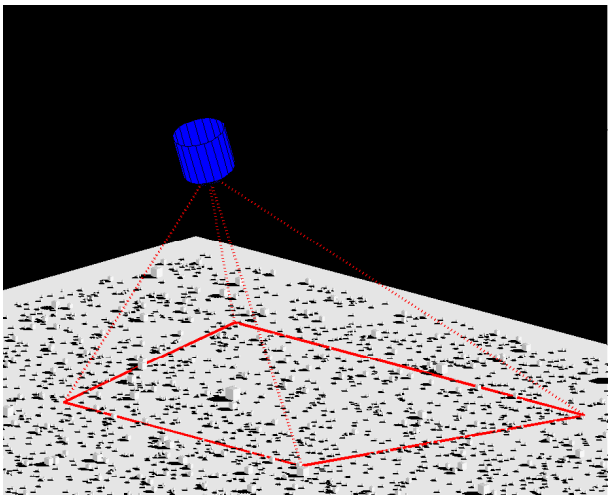

Figure 8. Simulated lander performing synthetic imaging during descent

Hazard Field Generation—The simulation environment can generate random rock fields of varying densities based on the Golombek model [18] for the fractional area of a region which is covered by rocks. While the model was originally developed for Mars, studies of rock density on the Moon have shown that the model remains valid and that rock density varies between 3% and 18% [21]. For convenience in generating images with shadow fields, simulated rocks are rectangular prisms. This assumption leads to greater ambiguity in finding the center point of the rock in the processing step, but as the technique remains effective on rocks of this type, the model is taken as conservative at the current level of development.

Synthetic Imaging—Synthetic images are created within the simulation environment and used in the estimation system as measurements. The images are created using the OpenGL renderer in MATLAB™ by tuning the rendering properties to match those of the simulated descent camera. All images are rendered in greyscale at a constant size of 640 × 480 pixels. An example image with blur and noise is shown in Figure 9. This image was taken over a randomly generated rock field with 7% coverage at an altitude of $30m$. Rock heights vary between $0.25m$ and $2.5m$, with frequencies dictated by the Golombek exponential size distribution.

Experimental Setup—In order to evaluate the benefits of the MPC trajectory modification scheme, a set of experiments was devised using the simulation environment: various trajectories were flown over the same hazard field and the uncertainty of the occupancy grids created were compared. The trajectories all had the same initial state and the same target landing site. While hazard maps were created, no landing site redesignation was performed during the descent. For every set of trajectories, the same rock field was imaged and the on-board mapping algorithm was initialized with the same a

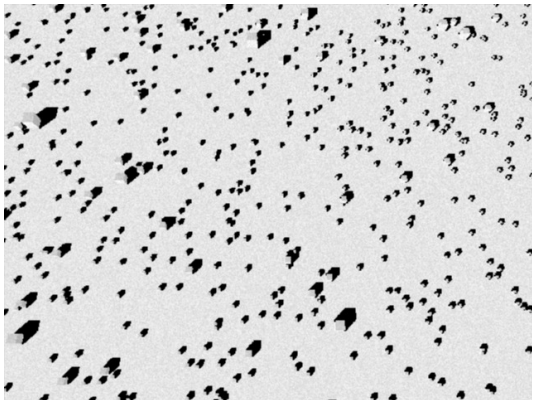

Figure 9. Example synthetic image

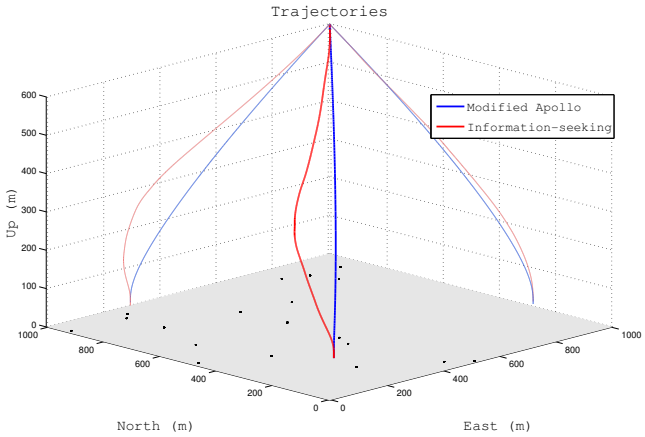

Figure 10. Trajectory comparison

priori knowledge. The filtering and mapping algorithms were identical for each run; the guidance algorithm was the only difference between runs.

6. RESULTS

The quality of hazard maps generated using the nominal and modified guidance schemes are compared to verify that the guidance system is driving the hazard map uncertainty towards zero. The shape of the modified trajectory is explored, as well as imaging and pose metrics which show the increases in information which explain the improvement in map quality. The various hazard maps are then considered in conjunction with spacecraft touchdown delivery errors in order to create a total probability of landing success, and the information-seeking guidance is shown to improve the availability of landable terrain.

Trajectory Modification

The guidance algorithms were flown over a sparse map, only 2.5% rock coverage, in order to establish baseline performance. The sun is modeled to be behind the lander during the descent, so the γ angle between the lander and the landing zone starts at greater than $90°$. The first trajectory was generated by the baseline MAG algorithm, which considers only fuel consumption and time-to-go constraints; no information-seeking augmentation inputs were added. This trajectory is shown in blue in Figure 10. The adjusted trajectory, generated using the MPC scheme and the cost function given in Equation 11 in order to drive down map uncertainty. The lander was initialized using the same position and velocity state; the trajectory generated is shown in red.

While both trajectories succeed in bringing the lander to the target site, the information-seeking trajectory shows more loft, hence more opportunities for imaging the target site. The improvement in imaging coverage is measured by examining the number of grid points in the target area which are imaged at a given altitude. The maximum number of points is 1681, as shown in Figure 11, with altitude decreasing to the left.

As seen in blue, the baseline MAG trajectory does not image any portion of the target area until almost $110m$ altitude, which allows for limited convergence of the hazard filter before a target site must be designated. The information-seeking guidance adjusts the trajectory to image the landing area at high altitude, $600m$, which enables any large hazards to be imaged if present. The effect of the MPC N-

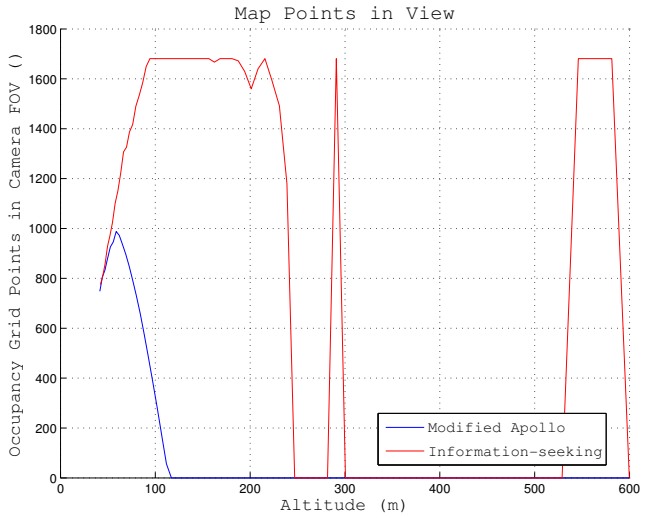

Figure 11. Site imaging coverage comparison

step horizon then becomes apparent: rather than continuing to pursue imaging at high altitude where image resolution is limited, the MPC scheme expends control effort setting up imaging opportunities at lower altitudes, where higher image resolutions provide much larger information returns. The MPC scheme provides imaging at almost the maximum possible between $225m$ and $95m$ altitude; the drop off in grid points imaged below $100m$ is due to the limited size of the camera field of view.

While the imaging coverage at low altitude achieved by the modified trajectory is obviously much greater, a few individual controller actions are examined to further expose the effect of the long-horizon prediction. In all cases, the size of the information-seeking adjustment is relatively small: no single control action, particularly at high altitude, can provide a large drop in map entropy. Rather, by using the model-predictive framework, small control actions can be accumulated to provide outsize improvements at low altitude. To illustrate this idea, two acceleration adjustment options, and the trajectories which result, available to the information-seeking scheme at high altitude are shown in Figure 12. The green area represents the landing zone for which the hazard occupancy grid is created in all cases. The blue quadrilaterals represent the predicted image field of view

along the proposed trajectory, up to the cutoff altitude, while the black shows the path which the lander traces, projected onto the ground plane.

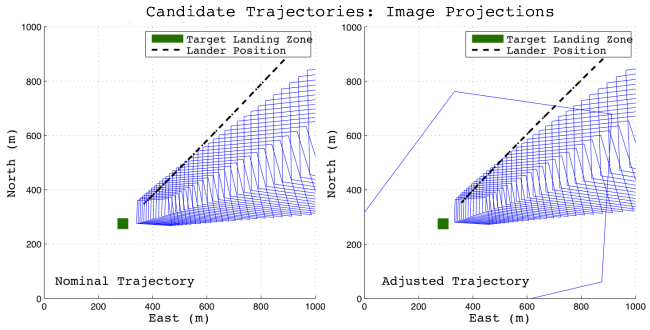

Figure 12. Image projections from candidate trajectories computed at $600m$

The trace on the left shows the unadjusted coverage up to a controller cutoff altitude of $150m$. This represents the images captured by the nominal trajectory, which captures no images of the landing zone during this altitude region. The trace on the right shows the adjustment made to the trajectory, which enables the spacecraft to take just one image of the target site. While still an improvement over the nominal case, a single large hazard is imaged and initialized into the hazard map, this image is taken at high altitude and does not enable any of the low-altitude imaging shown in Figure 11.

However, these incremental changes in trajectory and imaging opportunities build up over time, which the MPC scheme can leverage even when none of the acceleration adjustments within its capability can enable the landing zone to be imaged at the next opportunity. Following this scheme, by $320m$ altitude, all the intermediate trajectory adjustments have caused bulk movement of the trajectory bundle, as represented by the same image coverage plots shown in Figure 13. Both plots show the trajectory created by the information-seeking method up to this altitude, as seen in gray.

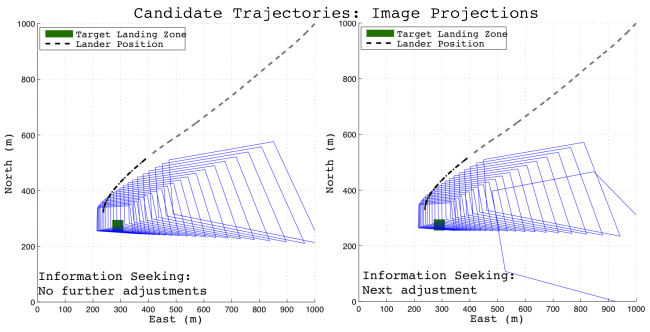

Figure 13. Image projections from candidate trajectories computed at $320m$

The trace on the left shows the trajectory which would be followed from $320m$ to the cutoff altitude if no further adjustments were made of the acceleration profile. Even without further refinement, there is already good image coverage achieved at information-rich altitudes. However, information-seeking adjustments continue, and the trajectory predicted by control choice actually taken by the spacecraft is shown on the right. This predicted trajectory, resulting from the choice made at $320m$, actually achieves less image coverage. However, given that the entire trajectory bundle enables similar, high coverage, the MPC algorithm discriminates the trajectory choices based more on the predicted entropy, which is influenced by the camera view angle. The trajectory adjustments made at high altitudes have put the lander in a state which is rich with imaging potential; consequently, at these lower altitudes, the MPC algorithm can fine tune the trajectory to minimize map uncertainty based on the terrain pattern which has emerged in the estimation system.

At even lower altitude, about $100m$, the modified trajectory also overflies the target site by $22m$ in the illumination direction, and so is able to image the target site from the shadowed side and provide un-occluded measurements of rock size. The angular measure of view angle γ shows that relative to the center of the landing area, the MPC scheme positions the lander with γ less than 90° at an altitude of $110m$, with variations along the landing area. The view angle is shown below in Figure 14.

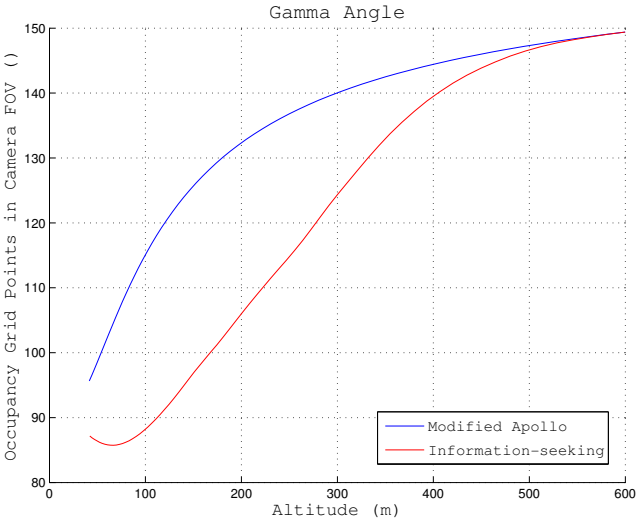

Figure 14. Gamma angle comparison

The occupancy grid uncertainty in the maps created during these two descents is also considered by looking at the total map entropy, shown in Figure 15.

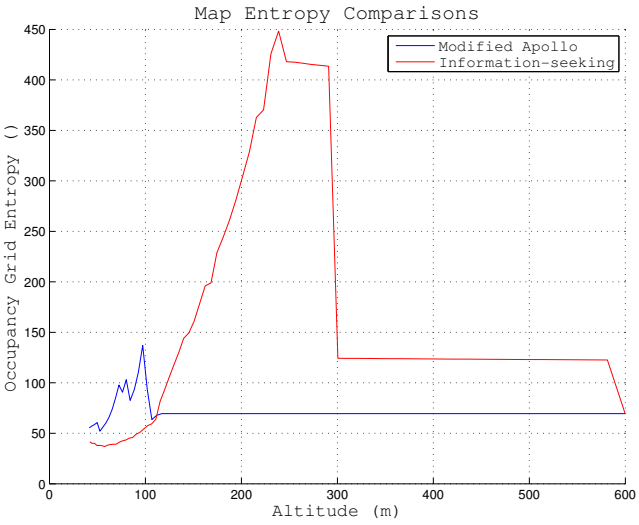

Figure 15. Map entropy comparison

The initial map entropy, on the right, starts at the same level and at the evaluation altitude, $40m$, the information-seeking trajectory shows that total map entropy is lower. Given the very low rock density of the hazard field at the landing area, the prior probability of occupation is also very low. The resulting prior map entropy is, therefore, extremely low. As such, when hazard objects are detected and added to the map, the overall entropy can actually *increase* due to the uncertainty associated with the object, particularly at high altitude when information is poor. While the entropy measure increases, the overall system understanding of the hazard field layout has been improved. The issue with increasing entropy with measurements in very low probability event systems has been documented and an alternative approach which seeks to maximize the *change* in system entropy has been developed [22], and is under consideration for future use in hazard mapping.

The entropy increase in the information-seeking trajectory map after the second image taken by the lander during the information-seeking descent is due to a large rock detected and added to the map; the large filter covariance for this object implies a large region of the map in which the system entropy has increased, as explained above. The large increase at $300m$ altitude corresponds to an image measurement of the whole area at enough resolution that small hazards are detected and added to the map. Further measurements drive the system uncertainty down as the filtered objects converge to the true size and layout, as evidenced by the following decrease.

The map entropy generated by the MAG baseline guidance is constant at the level of the prior until $110m$ altitude, as there are no image measurements of the target site collected until this low altitude. There is a small rise in the entropy as objects are added to the map, but it is not nearly as large as that in the information-seeking trajectory as only a portion of the map is imaged and so the number of objects added at once is much smaller than in the adjusted map, and the low altitude of the measurements implies high image resolution and hence less uncertainty in the measurements.

The cost of the information collected through trajectory adjustments is measured in increased propellant usage. The MPC-generated trajectory consumes and additional $1kg$ of propellant, representing a 14% increase in fuel usage over the nominal MAG trajectory.

Map Uncertainty and Landing Success

While entropy is a measure of map convergence and an effective variable for use in optimization, direct measures of the probability map uncertainty and terrain classification can be more illustrative of overall system performance. Low map uncertainty should lead to tight contours of high probability around hazard objects and areas of very low probability in regions which are measured as clear. A map with these low uncertainty characteristics should allow better landing site selection, as more points can be labeled with a clearly identified risk level due to occupancy, and overall risk due to a map uncertainty should be low.

When evaluating maps for landing site selection, the realities of lander control are captured by incorporating a probability distribution for lander delivery error relative to the desired touchdown location. This error is modeled using a two-dimensional Gaussian over a $2m$ square area, and is shown in Figure 16.

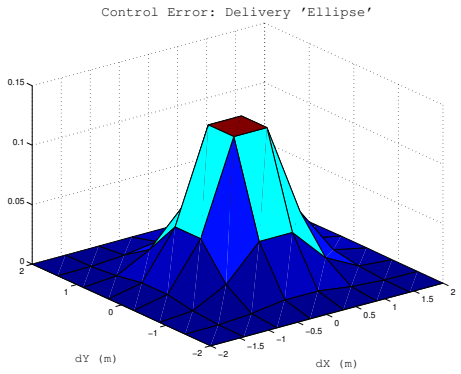

Figure 16. Delivery error distribution

The probability of landing success is evaluated at each grid point, assuming that point was selected as the final touchdown site, using combination of the estimated hazard probability map and the delivery uncertainty. A threshold is applied to this probability: if the lander has more than a 1% chance of failure due to either an estimated hazard being too near the target site or the ambiguity of the hazard layout being too large for a confident assessment of the terrain, the site is deemed unlandable. The delivery Gaussian kernel is convolved with the occupancy grid probability map in order to make the assessment of every point in the hazard map. The results of this convolution are shown in Figure 17. Landing success rates using true hazard field distribution, implying an occupancy grid composed only of the values 0 and 1, are shown on the left, with the estimated distributions shown on the right. Points where the probability of a successful landing is calculated to be higher than 99% are shown in green; points shown in red have an unacceptably high chance of failure.

Both estimated maps show the same basic hazard pattern as the true map, which is taken as confirmation that the estimation and mapping algorithms are performing well and that the image processing technique is capable of providing the necessary information to the filter. However, the map produced by the MAG trajectory shows a large swath of terrain as unlandable due to high uncertainty, as the lander was not able to image it during descent. This map contains no information on the actual distribution of hazards in this area. On this trajectory, the lander would thus have a limited choice of landing sites for selection due to a lack of information collected during the descent. The map produced by the information-seeking trajectory shows convergence to the true distribution over the whole map area. Two additional hazard objects can be seen at the bottom of this map which are not in the true hazard field shown on the left. These objects are being estimated in the EKF but are centered slightly outside the boundaries of the occupancy grid maps shown here. The uncertainty associated with them causes small regions of the occupancy grid to be labeled as unsuitable.

To quantify the improvement in mapping, the true proportion of landable sites is compared to the estimated proportion. The rock field is distributed such that 93% of the terrain is safe to land on, given a perfect map. In contrast, the MAG map identifies only 58% of the terrain as landable, which severely limits the terrain available for site selection. The information-seeking trajectory produces a map with 91% of the area identified as landable, which is almost the maximum possible. The improved map identifies an additional 33% of

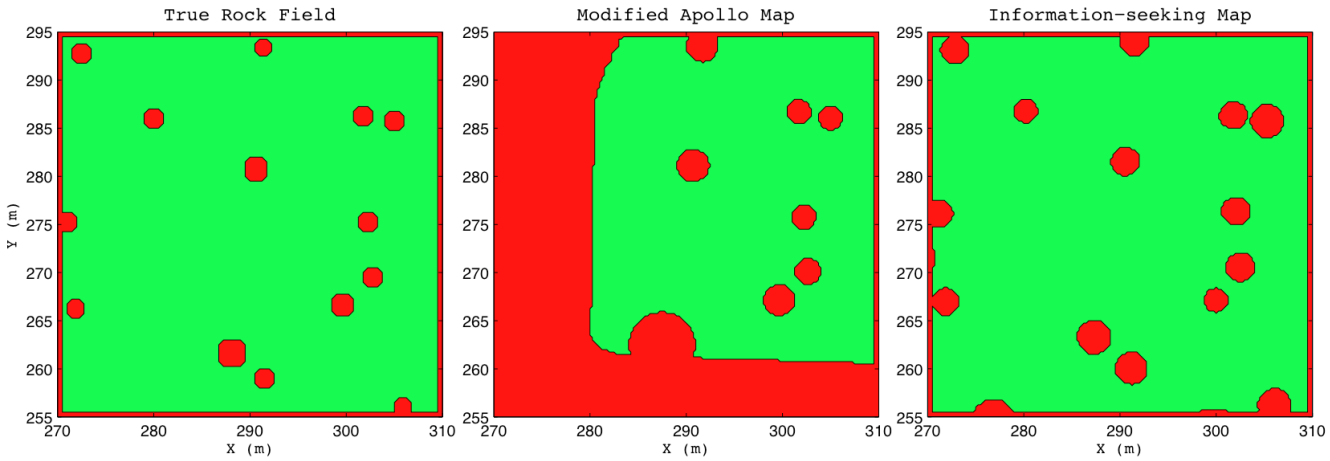

Figure 17. Classified terrain maps: probability of success

the target zone as landable. In addition, the terrain classifications produced by the lander during the actively-modified descent is 97% accurate in its identification when compared to the true rock field, which confirms that occupancy grid has converged to the true rock field.

7. Conclusion

A method for systematically producing information-seeking trajectories for hazard mapping during lunar landing has been presented. The use of a heritage spacecraft guidance algorithm ensures that the lander can make a soft landing on the surface, while the augmentation commands, which are generated using an model-predictive controller which aims to minimize hazard map uncertainty, improve the chances of a successful landing site redesignation at the decision altitude.

Simulation results have demonstrated that hazard map uncertainty can be reduced using the information-seeking system, at the cost of additional fuel usage during descent.

Acknowledgments

This research was funded in part by a Stanford Graduate Fellowship.

References

[1] C. Epp, E. Robertson, and T. Brady, "Autonomous landing and hazard avoidance technology (alhat)," in *Aerospace Conference, 2008 IEEE*, March 2008, pp. 1 –7.

[2] E. S. Crane and S. M. Rock, "Influence of trajectory on accuracy of hazard estimation during lunar landing," in *AIAA Guidance, Navigation, and Control Conference*, AIAA. Minneapolis, MN: AIAA, August 2012.

[3] T. Brady, E. Robertson, C. Epp, S. Paschall, and D. Zimpfer, "Hazard detection methods for lunar landing," in *Aerospace Conference, 2009 IEEE*, March 2009, pp. 1 –8.

[4] C. Epp, T. Smith, and H. NASA, "Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," in *2007 IEEE Aerospace Conference*, 2007, pp. 1–7.

[5] A. Johnson, J. Collier, and A. Wolf, "Lidar-based hazard avoidance for safe landing on mars," in *AAS/AIAA Space Flight Mechanics Meeting*, 2001, pp. 1091–1099.

[6] A. Johnson, A. Huertas, R. Werner, and J. Montgomery, "Analysis of on-board hazard detection and avoidance for safe lunar landing," in *Aerospace Conference, 2008 IEEE*, March 2008, pp. 1 –9.

[7] A. Huertas, Y. Cheng, and R. Madison, "Passive imaging based multi-cue hazard detection for spacecraft safe landing," in *Aerospace Conference, 2006 IEEE*, March 2006, p. 14 pp.

[8] L. Matthies, A. Huertas, Y. Cheng, and A. Johnson, "Stereo vision and shadow analysis for landing hazard detection," *IEEE International Conference on Robotics and Automation, ICRA 2008*, pp. 2735–2742.

[9] Y. Kim and J. Park, "Landing site searching algorithm of a quadrotor using depth map of stereo vision on unknown terrain," in *Infotech@Aerospace 2012*. American Institute of Aeronautics and Astronautics, Oct 2012.

[10] A. Howard and H. Seraji, "Multi-sensor terrain classification for safe spacecraft landing," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 40, no. 4, pp. 1122 – 1131, Oct 2004.

[11] N. Serrano, "A bayesian framework for landing site selection during autonomous spacecraft descent," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2006, October 9-15, 2006, Beijing, China*. IEEE, 2006, pp. 5112–5117.

[12] J. de Lafontaine, D. Neveu, and J.-F. Hamel, "Pseudo-doppler velocity navigation for lidar-based planetary exploration," in *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. American Institute of Aeronautics and Astronautics, Oct 2012.

[13] B. E. Cohanim and B. Collins, "Landing point designation algorithm for lunar landing," *Journal of Spacecraft and Rockets*, vol. 46, no. 4, pp. 858–864, July/August 2009, doi: 10.2514/1.42002.

[14] S. Kohlbrecher, A. Stumpf, and O. von Stryk, "Grid-based occupancy mapping and automatic gaze control for soccer playing humanoid robots," in *Proc. 6th Workshop on Humanoid Soccer Robots at the 2011 IEEE-*

RAS Int. Conf. on Humanoid Robots, Bled, Oct. 26th - Oct. 28th 2011.

[15] J. F. Seara, K. H. Strobl, and G. Schmidt, "Path-dependent gaze control for obstacle avoidance in vision guided humanoid walking," in *ICRA'03*, 2003, pp. 887–892.

[16] F. L. Markley, "Attitude error representations for kalman filtering," *Journal of Guidance Control and Dynamics*, vol. 26, no. 2, pp. 311–317, 2003.

[17] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46 –57, june 1989.

[18] M. P. Golombek *et al.*, "Rock size-frequency distributions on mars and implications for mars exploration rover landing safety and operations," *Journal of Geophysical Research*, vol. 108, no. E12, 2003.

[19] A. R. and Klumpp, "Apollo lunar descent guidance," *Automatica*, vol. 10, no. 2, pp. 133 – 146, 1974.

[20] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, ser. Intelligent Robotics and Autonomous Agents. Mit Press, 2005.

[21] "Rock distribution on the lunar surface," https://www.lpi.usra.edu/wiki/lunaref/index.php/Rock_Distributions_on_the_Lunar_Surface, Sept 2007.

[22] Z. Saigol, R. Dearden, J. Wyatt, and B. Murton, "Belief change maximisation for hydrothermal vent hunting using occupancy grids," in *Proceedings of the Eleventh Conference Towards Autonomous Robotic Systems (TAROS-10)*, 2010, pp. 247–254.

BIOGRAPHY

Eleanor S. Crane is a PhD candidate at Stanford University in the department of Aeronautics and Astronautics; her research is focused on hazard avoidance for planetary landing. She also works in this area as a GNC and Software Engineer at Moon Express, Inc, a lunar transportation and data services company which is also a leading contender for the Google Lunar X-prize. Her background includes work as a flight software engineer for the Lunar Atmosphere and Dust Environment Explorer (LADEE) mission at NASA Ames Research Center, in the area of state estimation. Eleanor received her Masters degree from Stanford in 2007, and her Bachelors degree from MIT in 2006.

Stephen M. Rock is a professor in the Department of Aeronautics and Astronautics and is the director of the Aerospace Robotics Laboratory at Stanford University. Dr. Rock's interests include the development and experimental verification of advanced control techniques for robotic and autonomous vehicle systems. Dr. Rock joined the Stanford faculty in 1988. Prior to that, he led the Controls and Instrumentation Department of Systems Control Technology, Inc. Dr. Rock is a Fellow of the AIAA.