Belief Space Planning for Reducing Terrain Relative Localization Uncertainty in Noisy Elevation Maps

Eugene Fang, P. Michael Furlong, and William "Red" Whittaker

***Abstract*— Accurate global localization is essential for planetary rovers to reach mission goals and mitigate operational risk. For initial exploration missions, it is inappropriate to deploy GPS or build other infrastructure for navigating. One way of determining global position is to use terrain relative navigation (TRN). TRN compares planetary rover-perspective images and 3D models to existing satellite orbital imagery and digital elevation models (DEMs) for absolute positioning. However, TRN is limited by the quality of orbital data and the presence and uniqueness of terrain features. This work presents a novel combination of belief space planning with terrain relative navigation. Additionally, we introduce a new method for increasing the robustness of belief space planning to noisy map data. The new algorithm provides a statistically significant reduction in localization uncertainty when tested on elevation data produced from lunar orbital imagery.**

I. Introduction

Accurate absolute localization is a vital component of planetary exploration missions. Vehicles operating on the moon that cannot survive lunar night are constrained to fourteen days of operations, which limits permissible margins of error when navigating to mission goal locations. Missions which track time-sensitive phenomena, e.g. capturing specific views under dynamic lighting conditions, or solar-powered robots chasing insolation and avoiding shadows [19][1][5], depend on accurate localization. Many craters and canyons look alike from a local perspective (Fig. 1); precise absolute localization alleviates confusion in position that can end a mission.

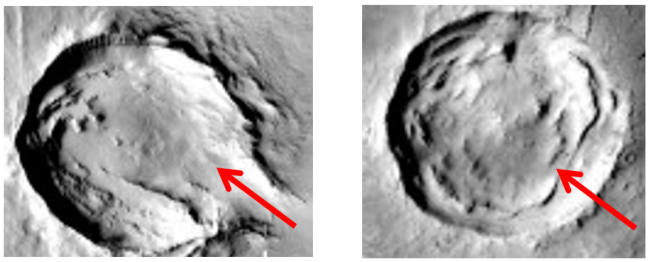

(a) Safe entry point into crater (b) No safe entry point into crater

Fig. 1: Entering two craters along the red arrows may appear similar when viewed from a rover's cameras. Accurate localization can ensure that a rover enters the correct crater (1a) through a safe entry point. Inaccurate localization can result in entering the wrong crater, causing mission-ending damage (1b).

Building localization infrastructure for new planetary bodies could be cost prohibitive, or be inappropriate in the case of initial exploration missions. Absent localization infrastructure, one can localize relative to terrain features. Terrain Relative Navigation (TRN) estimates robots' global positions by correlating egocentric data with an *a priori* map [17][2] produced from orbital data via long-baseline stereo or photogrammetry [15].

With TRN, localization accuracy depends on the presence and uniqueness of terrain features; the path taken by a rover affects its certainty of localization. Thus, it is natural to apply belief space planning techniques to improve the localization performance of the rover along the trajectory. This poses two interesting problems: First, where other approaches rely on discrete beacons [14], belief space planning in terrain relative navigation presents, to the resolution of the orbital map, a continuum of beacons, which may not all be unique. This can cause planned paths to wander circuitously in order to gain more information from TRN before heading towards the goal. Second, compounding this problem, is the fact that orbital maps are noisy and so the extracted DEMs may not be accurate representations of the terrain geometry that rovers will experience on the surface. A naive approach to belief space planning may be fooled by what appear to be unique landmarks on the terrain map, but are in fact artifacts of sensor noise.

To address the first problem, we modified the belief roadmap algorithm [14] by adding a path length penalty term to the optimization criterion. To address the second problem we produced the innovative strategy of sampling noisy versions the terrain maps in order to find landmarks that are less likely to be mirages due to inaccuracies in the constructed DEMs. We show, on data of the lunar surface, that our approach reduces terminal location uncertainty over a belief space planner that does not use noisy map sampling, with statistical significance and moderate effect size.

The benefit of our innovation can be viewed in two lights. Primarily, it can improve the localization performance of surface vehicles which are localizing in noisy maps. Alternatively, the algorithm can be viewed as improving the operational tolerance for uncertainty in preliminary mapping, reducing demands on the fidelity of precursor sensors, potentially reducing mission cost.

II. Background

A. Terrain Relative Navigation

The most recent Mars rover, the Mars Science Laboratory (MSL), corrects for odometry drift by relying on humans on Earth to match terrain features in rover imagery to the

same features spotted in orbital images. This manual correlation is not suitable for high-cadence autonomous rovers. Automated localization systems have been explored, with techniques ranging from rover-orbital image matching [7], horizon matching to DEMs [4], and local 3D matching to DEMs [12].

SLAM is often used in robotic applications to correct for drift in localization when there are sufficient loop closures [16][6][11]. However, SLAM is unnecessary due to already having prior satellite maps and does not solve the problem of localizing and reaching mission goals within these maps.

B. Belief Space Planning

The performance of any approach to terrain relative navigation fundamentally depends on the quality of the data which produce the terrain model (e.g. orbital imagery and DEMs), the presence of terrain features, and the uniqueness of those features. Current global planners [10][3] for planetary rover exploration primarily plan to minimize traverse cost (distance, energy, etc.). These paths may take a rover into regions with sparse or ambiguous terrain features, resulting in high localization uncertainty. Belief space planning can produce paths that constrain the rover's localization uncertainty to within acceptable values.

Belief space planning is the problem of planning with imperfect state information, considering the predicted uncertainty of future states. This is a challenging problem since the dimensionality of the belief state space is much larger than the original state space. Prentice and Roy [14] tackle this problem with the Belief Roadmap (BRM) algorithm, which reduces the search space with a Probabilistic Roadmap (PRM) [9] and combines multiple extended Kalman filter (EKF) updates into a single linear transfer function for efficient calculation of posterior belief during planning. This methods assumes the existence of a finite set of beacons in the environment that are themselves accurately localized, whereas terrain relative navigation is performed on a continuum of "beacons," with each point having its own utility with respect to localization. Other approaches [13][18] incorporate variance into an augmented state representation to find a locally optimal control policy using a linear-quadratic Gaussian (LQG) framework. These approaches can only find a locally optimal policy which is highly dependent on the initialization, and stipulate that observation models and cost functions need to be sufficiently smooth. While these assumptions for smoothness may be valid in many beacon-based navigation scenarios, localization uncertainty from TRN can be discontinuous if a unique terrain feature suddenly appears or disappears from view. Other research investigates the problem of planning under uncertainty where there is no prior knowledge of the environment [8]. Our approach does not address this problem, and instead assumes the existence of prior orbital data.

III. Noisy-DEM-Tolerant Belief Space Planning

Given a map, start and goal positions, uncertainty in starting position, a model of localization uncertainty growth from robot odometry, and a particular method for terrain relative navigation, our objective is to plan a hazard-free path from start to goal that minimizes some function of the localization uncertainty at the goal. Additionally, it is desirable that this path is minimally affected by noise in the map that TRN uses for localization. This work uses a simple implementation of horizon matching as the means for terrain relative navigation as an example, but is agnostic of the particular TRN implementation.

The developed method partitions planning into two phases:

1) **Pre-processing phase:** Prediction of TRN uncertainty, noisy DEM sampling, and BRM graph construction
2) **Query phase:** Graph search through the BRM given start and goal positions

A. Prediction of Terrain Relative Navigation Uncertainty

To predict the horizon matching TRN uncertainty at a particular position (described by a 2D covariance), the similarity between two horizons must first be defined and calculated. The horizon at a position in the DEM is rendered by ray-tracing 360-degree "panoramas" at 1-degree intervals (Fig. 2). The similarity between two horizons is defined as an exponential of the sum of squared differences between them (described in [4]), where h^1 and h^2 are the panoramic elevations of two horizons.

$$similarity = \frac{1}{\sqrt{2\pi}} \exp\left(- \frac{\sum_{n=1}^{360} (h_n^1 - h_n^2)^2}{2} \right)$$

We assume that a means for estimating orientation (e.g. sun compass or star tracker) is available, so no rotational alignment of panoramas is necessary.

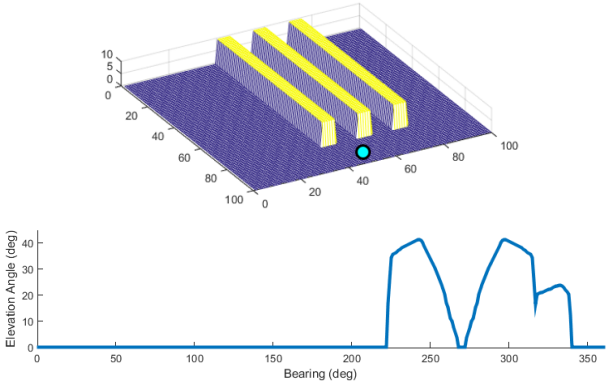

Fig. 2: Example of a rendered horizon (bottom) at the cyan point in a contrived environment (top).

Computing the similarity between a given position's horizon to all other positions' horizons yields a positional uncertainty probability distribution of TRN at that given position (Fig. 3), which is then fit with a 2D covariance centered at that position. The 2D covariance is only an approximation of the underlying uncertainty which may actually be multimodal and non-continuous, but enables efficient estimation of uncertainty using the extended Kalman filter when planning.

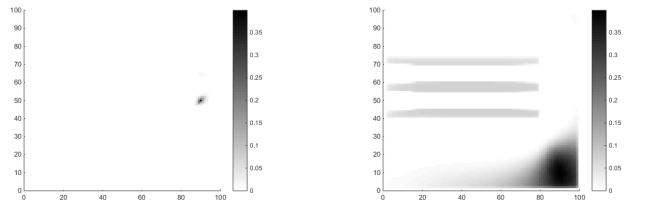

(a) The horizon at position (90, 50) is less similar to other horizons. This results in a positional uncertainty with low variance.

(b) The horizon at position (90, 10) is similar to many other horizons. This results in a positional uncertainty with high variance.

Fig. 3: The positional uncertainty of TRN due to horizon similarity at different positions in a contrived environment. Darker = more similar; Lighter = less similar.

The process is repeated for every cell in the DEM to obtain a prediction of TRN uncertainty at every cell in the map (Algorithm 1) and can be visualized by plotting the trace of the covariance at every cell (Fig. 4).

Algorithm 1 Algorithm for computing TRN uncertainty covariances at every position in the DEM

Init $covars$
for $pos_1 \in DEM$ **do**
  Init $probDist$
  $h_1 \leftarrow getHorizonAt(pos_1, DEM)$
  for $pos_2 \in DEM$ **do**
    $h_2 \leftarrow getHorizonAt(pos_2, DEM)$
    $probDist[pos_2] \leftarrow computeSimilarity(h_1, h_2)$
  end for
  $covars[pos_1] \leftarrow computeCovar(probDist)$
end for

B. Multiple DEM Sampling to Mitigate Noise in DEMs

The described method for estimating the positional uncertainty when performing TRN at a given position is subject to noise in digital elevation models; a seemingly unique horizon may only be unique due to noise in the DEM. The effect of these "false positive" low uncertainty regions is lessened by sampling multiple noisy DEMs from the original noisy DEM, generating multiple probability distributions from these noisy DEMs, and then taking the positional maximum across all probability distributions before computing covariances (Algorithm 2). As the DEM noise level increases, more sample iterations are needed to adequately counteract the noise. In most of this work, five iterations of noisy samples are used.

The effect of noisy DEM sampling is that positions which are unique *despite the presence of noise* are distinguished from positions which seem unique *only because of noise* (Fig. 5).

C. Belief Roadmap Graph Construction

To reduce the search space, a probabilistic roadmap graph [9] is constructed. The graph is made up of nodes at

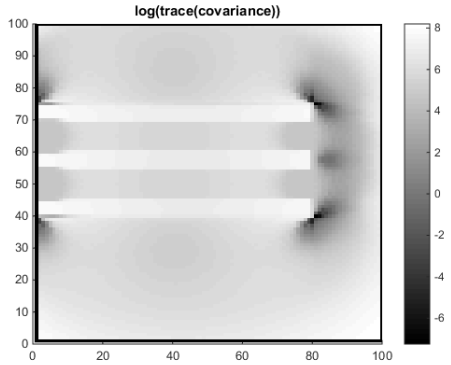

Fig. 4: A plot of the TRN uncertainty at every position over the contrived environment. Darker areas have lower uncertainty, and should be favored during planning.

Algorithm 2 Modification of Algorithm 1 to mitigate noise in DEMs

Init $covars$
for $pos_1 \in DEM$ **do**
  Init $probDist_{max}$
  for $i \in numIters$ **do**
    Init $probDist_i$
    $noisyDEM \leftarrow generateNoisyDEM(DEM)$
    $h_1 \leftarrow getHorizonAt(pos_1, noisyDEM)$
    for $pos_2 \in DEM$ **do**
      $h_2 \leftarrow getHorizonAt(pos_2, noisyDEM)$
      $probDist_i[pos_2] \leftarrow computeSimilarity(h_1, h_2)$
    end for
    $probDist_{max} \leftarrow max(probDist_{max}, probDist_i)$
  end for
  $covars[pos_1] \leftarrow computeCovar(probDist_{max})$
end for

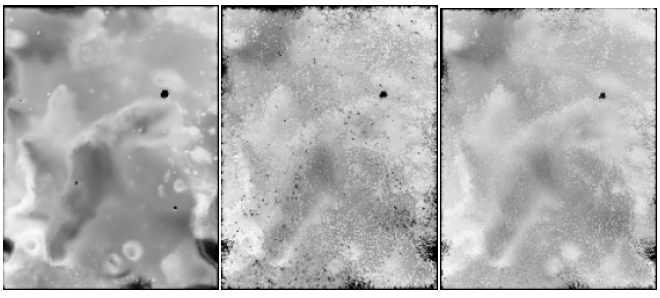

(a) TRN uncertainty from ground truth DEM

(b) TRN uncertainty from noisy DEM

(c) TRN uncertainty from 5x noisy sampled DEM

Fig. 5: TRN uncertainty plotted as $log(tr(covariance))$ of a lunar DEM. Darker shades indicate lower uncertainty, while whiter shades indicate higher uncertainty. Dark points of seemingly low uncertainty in (5b) are false positives and aren't actually present in (5a). Sampling the noisy DEM reduces these false positives (5c).

randomly sampled positions and hazard-free edges between nodes within a certain distance of each other. Hazards are defined by cells in the elevation map where the slope exceeds a pre-defined limit. The slope of a cell $s(x, y)$ is computed as the average of the slopes for the 4-connected neighbourhood of the cell (x, y). Where the slope between any two neighbours is computed $|M(x, y) - M(x', y')|/W$, where $M(\cdot)$ is the elevation map, and W is the width of a cell in the DEM and $(x', y') \in \{(x-1, y), (x+1, y), (x, y-1), (x, y+1)\}$.

Multiple extended Kalman filter updates for each edge (consisting of many odometry and terrain relative navigation measurements) are precomputed into one-step transfer functions for the covariance as described in the original the belief roadmap (BRM) paper [14]. This enables more efficient search later on, where an EKF is used to propagate uncertainty belief.

D. Graph Search Using a Modified BRM Algorithm

To find the lowest cost path from start to goal, the belief roadmap algorithm keeps track of the covariance, best path, and lowest cost to each node. Breadth first search is performed, initializing the search queue with the start location and covariance. The covariance for each neighboring node is computed by applying the precomputed one-step transfer functions, and neighbors are only added to the search queue if the cost to that node is less than the lowest previously seen cost for that node. Search continues until the queue is exhausted, at which point the lowest cost path to the goal will have been found if it exists.

The original BRM algorithm defines the cost function to be the trace of the goal covariance, $tr(\Sigma_{goal})$. While this works in a domain with a limited number of discrete beacons for localization, it becomes problematic in the TRN domain. Here, localization can be performed continuously everywhere and there exist points with very low localization uncertainty. Using the original cost function, planned paths may wander to reach lower uncertainty locations to maximize information gain from TRN before heading towards the goal. To drive the path towards the goal, a new cost is defined as $pathLength* \alpha + tr(\Sigma_{goal}) * \beta$. The effect of weighted cost can be seen in Fig. 6.

An example of a path planned by the algorithm is shown in Fig. 7. Here, the initial starting position is unknown. The shortest path cannot differentiate between the two valleys, resulting in a high final uncertainty. The belief-optimal path exits the valley for more accurate TRN localization to decrease uncertainty before heading towards the goal.

IV. Experiment

A. Setup

To evaluate the effectiveness of our approach for planning low-final-uncertainty paths when performing TRN, a Monte Carlo study was performed on a DEM near the Nobile crater region of the Moon (Fig. 8). The DEM covers a 20x15 km area with a resolution of 100 meters per pixel and has $\sim$ 1.6 km elevation range. To account for the effect of noise and errors in the DEM, planning is performed on a

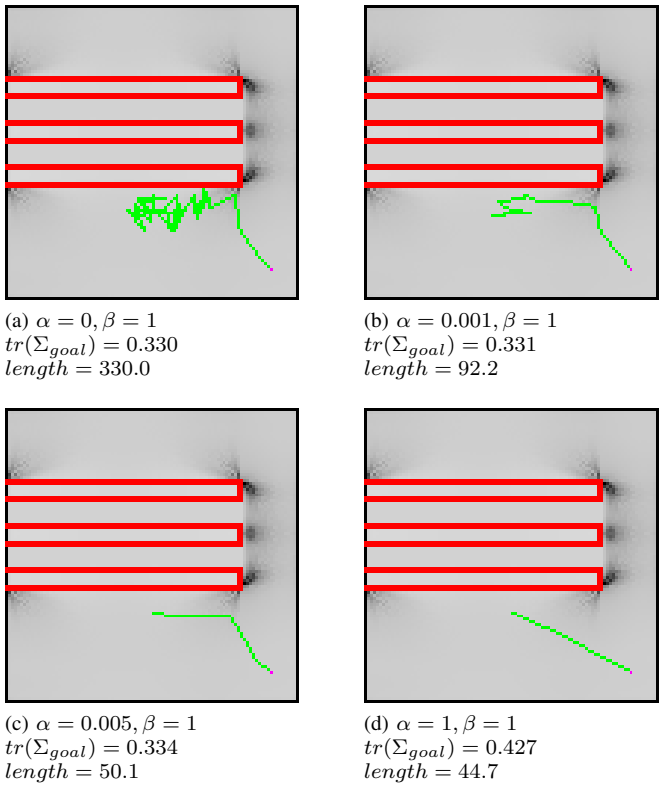

(a) $\alpha = 0, \beta = 1$
$tr(\Sigma_{goal}) = 0.330$
$length = 330.0$

(b) $\alpha = 0.001, \beta = 1$
$tr(\Sigma_{goal}) = 0.331$
$length = 92.2$

(c) $\alpha = 0.005, \beta = 1$
$tr(\Sigma_{goal}) = 0.334$
$length = 50.1$

(d) $\alpha = 1, \beta = 1$
$tr(\Sigma_{goal}) = 0.427$
$length = 44.7$

Fig. 6: Using the original unweighted BRM cost function (6a), the path takes circuitous paths that maximize information gain from TRN before heading towards the goal (magenta). Increasing the weighted cost of path length (6b, 6c) trades localization uncertainty for shorter paths. As α is increased, the belief-state planner converges to the shortest path planner (6d). Red loci are slope hazards.

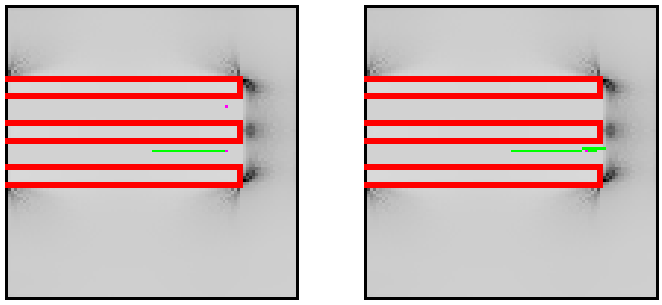

(a) The shortest path to the goal with cannot differentiate between the two valleys, yielding a high final uncertainty; the robot could be at either magenta point.

(b) The belief-optimal path leaves the valley to drastically lower uncertainty before heading back towards the goal; the robot knows it is at the magenta point.

Fig. 7: A comparison between shortest (left) and belief-optimal (right) paths when the starting position is unknown. Uncertainty is lowered ($tr(\Sigma_{goal}) = 4.09$ vs. 0.40) at the cost of path length ($length = 25.0$ vs. 37.1). Red loci are slope hazards.

version of the DEM with added Gaussian noise on elevation, while evaluation is done using the original unaltered DEM serving as "ground truth" to render rover-perspective horizons. The roadmap used for planning has 10,000 nodes and connectivity between nodes $\leq$ 1 km apart. Areas with slope greater than 15 degrees are considered untraversable hazards. Odometry error is set at $\sigma_{odom} = 10\%$ of traversed distance. α and β in the cost function are experimentally chosen to be 0.05 and 1 respectively to qualitatively balance path length and uncertainty. TRN is performed every 100 meters of traverse. The uncertainty at the start of each traverse is zero.

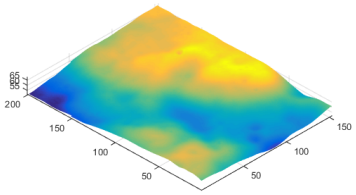

Fig. 8: A 20x15 km lunar DEM near the Nobile crater region of the Moon with a $\sim$ 1.6 km elevation range.

For each pair of 50 random start and goal points and over a range of added elevation noise in the DEM ($0 \leq \sigma \leq 20$ meters), the shortest and belief-optimal paths are planned. The rover traverse for each path is then simulated, performing terrain relative navigation by matching the simulated rover-perspective horizon from the unaltered "ground truth" DEM to horizons rendered from the noisy DEM and using an extended Kalman filter for localization. The terminal localization uncertainty for each belief-optimal path is compared to that of the shortest path.

B. Results

Fig. 10a shows the difference in final localization uncertainty when planning shortest and belief-optimal paths across 50 trials. With no added noise to the DEM ($\sigma = 0m$), belief-optimal planning yields a statistically significant reduction in terminal pose uncertainty, with p $<$ 0.001 and a moderate effect size (d=0.671). On average, the final localization uncertainty is reduced by 21% at the cost of extending driving distance by only 13%. With increasing levels of DEM noise, reduction in final uncertainty loses its statistical significance; it is only statistically significant (p $<$ 0.01) while DEM noise $\sigma \leq 10m$. With 5x sampling (Fig. 10b), belief-optimal paths are more robust to the noise in the DEM when compared to baseline BRM planning, resulting in paths that are more likely to decrease localization uncertainty as noise levels rise. Here, reduction in uncertainty is statistically meaningful to higher levels of DEM noise, up to around $\sigma \leq 14m$ ($p < 0.001, d = 0.399$), compared to $\sigma \leq 10m$ ($p < 0.001, d = 0.447$) of the baseline algorithm.

The effect of the number of noisy samples when planning on a DEM with significant noise ($\sigma = 20m$) is illustrated in Fig. 9 and Table I. When the number of samples is small ($<$ 10), the performance of the resampling algorithm is statistically indistinguishable from the baseline BRM. However, with a sufficient number of samples ($>$10) we see a small to moderate improvement in the terminal pose uncertainty of the vehicle with statistical significance (p $<$ 0.01).

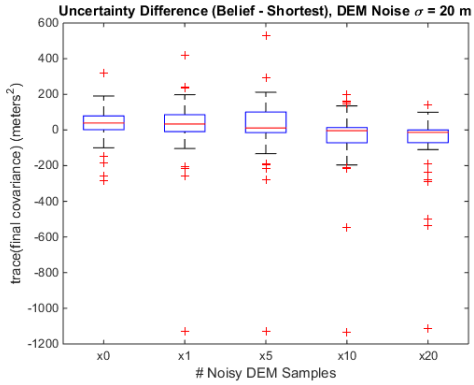

Fig. 9: At a given DEM noise level ($\sigma = 20m$), increasing the number of noisy DEM samples makes the effect of belief-optimal planning more statistically significant.

# Noisy samples	Average ($\pm$ std error) change in uncertainty ($meters^2$), n = 50	p	Cohen's d
0	29.3 (14.7)	N/A	N/A
1	14.5 (28.6)	0.561	0.092
5	9.0 (29.6)	0.455	0.123
10	-46.9 (28.3)	0.006	0.478
20	-76.1 (27.5)	$<$0.001	0.677

TABLE I: The result of paired T-tests comparing the effect of noisy DEM sampling to the baseline BRM algorithm (no noisy samples) at a given DEM noise level ($\sigma = 20m$). The baseline BRM algorithm does poorly, increasing terminal position uncertainty compared to shortest path planning. Increasing the number of noisy samples gives the desired result: a reduction in uncertainty even under extreme DEM noise.

Fig. 12 compares the paths planned from the same start and end points, with and without DEM noise. In the noise-free case, the baseline BRM path (with the added path length cost modifier) lowers the goal uncertainty at the cost of path length. Qualitatively, the path seems to favor the dark ridge of low uncertainty. When noise is introduced into the DEM, baseline BRM planning causes the path to deviate to seek out points of false-positive certainty that do not actually exist, resulting in both an increase of uncertainty and path length. When applying noise-tolerant belief-optimal planning with noisy DEM sampling, the uncertainty is once again lowered at the cost of path length. Qualitatively, the path more closely resembles the noise-free path, which suggests that noisy DEM sampling is effective at reducing the influence of DEM noise. The trade off between uncertainty and path length is visualized in Fig. 11 for all 50 random pairs of start and goal points.

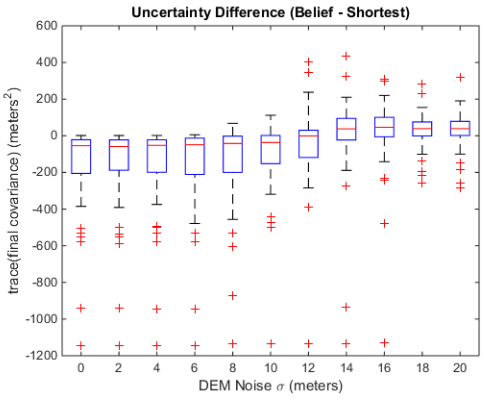

(a) Baseline BRM planning (with path length cost modifier)

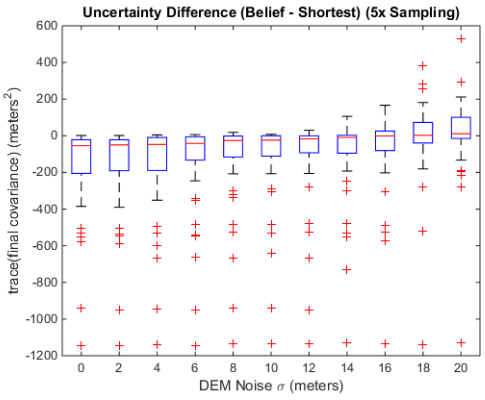

(b) DEM noise-tolerant planning (5x noisy DEM sampling)

Fig. 10: The difference in terminal position uncertainty (belief-optimal vs. shortest) as a function of DEM noise over 50 pairs of random start and goal points. Baseline BRM planning without noisy DEM sampling (10a) reduces final uncertainty with statistical significance while DEM noise $\sigma \leq 10m$. With 5x noisy DEM sampling (10b), belief-optimal planning is statistically effective to higher levels of DEM noise, up to around $\sigma \leq 14m$.

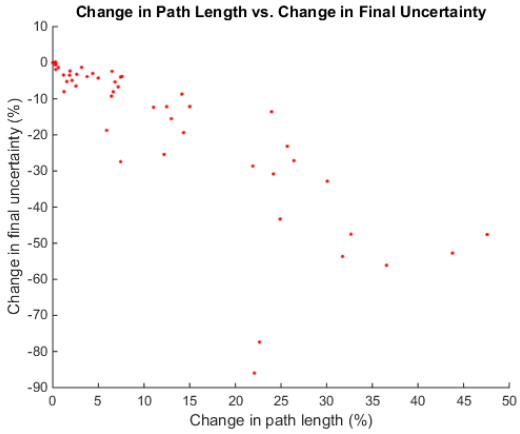

Fig. 11: Each point compares the % change in path length vs. % change in final uncertainty of belief-optimal planning compared to shortest-path planning in the case with no DEM noise for one of 50 trials. In general, the reduction in uncertainty is greater than the increase in path length.

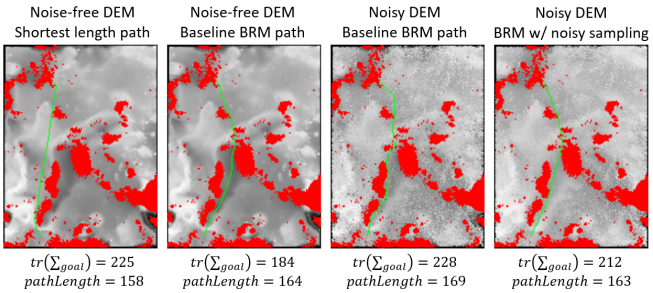

Fig. 12: A comparison of paths planned (green) from the same start and end points, with and without $\sigma = 12m$ of DEM noise. Darker areas have lower TRN uncertainty. Red areas are hazards.

V. Conclusions

We successfully applied belief space planning to terrain relative navigation for surface vehicles to reduce localization uncertainty. Further, we introduced a novel method for increasing a route's robustness to noise in digital elevation models through noisy DEM sampling. This reduces the likelihood of planned routes seeking out regions of seemingly low TRN uncertainty that do not actually exist. The methodology was evaluated on elevation maps of Lunar terrain through simulations over many traverses with varying DEM noise levels. In the case with no DEM noise, belief-optimal planning always reduced final localization uncertainty when compared to shortest path planning, with $p < 0.001$ and a moderate effect size (d=0.671). It decreased localization uncertainty by 21% at the cost of extending driving distance by only 13% when compared to shortest-path planning. Under the highest DEM noise levels tested ($\sigma = 20m$), the baseline BRM algorithm did poorly, *increasing* terminal position uncertainty by an average of $29.3m^2$ compared to shortest path planning. Noisy sampling greatly improved upon this, *decreasing* uncertainty by an average of $76.1m^2$ with statistical significance ($p < 0.001, d = 0.677$). Future work should characterize the number of noisy samples appropriate for different levels of DEM noise.

This work will enable planetary exploration rovers to reach science goals more accurately and decrease mission risk by compensating for fundamental limitations in terrain relative navigation and DEM quality. Characterizing performance over different terrain regions, more accurate noise models, and better capturing the multi-modal aspect of TRN are all more areas for future work.

Acknowledgements

This work was supported by the NASA Space Technology Research Fellowship under grant NNX16AM68H. We would like to thank Dr. Oleg Alexandrov for his assistance with the Lunar elevation maps.

REFERENCES

[1] D Andrews, Anthony Colaprete, Jacqueline Quinn, Donald Chavers, and Martin Picard. Introducing the resource prospector (RP) mission. In *AIAA Space 2014 Conference & Exhibition*, 2014.

[2] Patrick JF Carle, Paul T Furgale, and Timothy D Barfoot. Long-range rover localization by matching lidar scans to orbital elevation maps. *Journal of Field Robotics*, 27(3):344–370, 2010.

[3] Joseph Carsten, Arturo Rankin, Dave Ferguson, and Anthony Stentz. Global planning on the mars exploration rovers: Software integration and surface testing. *Journal of Field Robotics*, 26(4):337–357, 2009.

[4] Fabio Cozman, Eric Krotkov, and Carlos Guestrin. Outdoor visual position estimation for planetary rovers. *Autonomous Robots*, 9(2):135–150, 2000.

[5] Chris Cunningham, Joseph Amato, Heather L Jones, and William L Whittaker. Accelerating energy-aware spatiotemporal path planning for the lunar poles. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 4399–4406. IEEE, 2017.

[6] Jakob Engel, Thomas Schöps, and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *European Conference on Computer Vision*, pages 834–849. Springer, 2014.

[7] G Foil, C Cunningham, DS Wettergreen, and WL Whittaker. Onboard Detection and Correction of Orbital Registration Errors Using Rover Imagery. 2014.

[8] Vadim Indelman, Luca Carlone, and Frank Dellaert. Planning under uncertainty in the continuous domain: a generalized belief space approach. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 6763–6770. IEEE, 2014.

[9] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.

[10] P Chris Leger, Ashitey Trebi-Ollennu, John R Wright, Scott A Maxwell, Robert G Bonitz, Jeffrey J Biesiadecki, Frank R Hartman, Brian K Cooper, Eric T Baumgartner, and Mark W Maimone. Mars exploration rover surface operations: Driving spirit at gusev crater. In *Systems, Man and Cybernetics, 2005 IEEE International Conference on*, volume 2, pages 1815–1822. IEEE, 2005.

[11] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE Transactions on Robotics*, 31(5):1147–1163, 2015.

[12] AV Nefian, LJ Edwards, D Lees, L Keely, TJ Parker, and M Malin. Automatic Rover Localization in Orbital Maps. In *Lunar and Planetary Science Conference*, volume 48, 2017.

[13] Robert Platt Jr, Russ Tedrake, Leslie Kaelbling, and Tomas Lozano-Perez. Belief space planning assuming maximum likelihood observations. 2010.

[14] Samuel Prentice and Nicholas Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *The International Journal of Robotics Research*, 28(11-12):1448–1465, 2009.

[15] David E Shean, Oleg Alexandrov, Zachary M Moratto, Benjamin E Smith, Ian R Joughin, Claire Porter, and Paul Morin. An automated, open-source pipeline for mass production of digital elevation models (DEMs) from very-high-resolution commercial stereo satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 116:101–117, 2016.

[16] Robert Sim, Matt Griffin, Alex Shyr, and James J Little. Scalable real-time vision-based SLAM for planetary rovers. In *IEEE IROS Workshop on Robot Vision for Space Applications*, pages 16–21, 2005.

[17] Fridtjof Stein and Gerard Medioni. Map-based localization using the panoramic horizon. *IEEE Transactions on Robotics and Automation*, 11(6):892–896, 1995.

[18] Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278, 2012.

[19] David Wettergreen, Benjamin Shamah, Paul Tompkins, and William Whittaker. Robotic planetary exploration by sun-synchronous navigation. *i-SAIRAS, Montreal, Canada*, 2001.