

Covariance Based Terrain Mapping for Autonomous Mobile Robots

Lennart Werner¹, Pedro Proença¹, Andreas Nüchter², Roland Brockers¹

Abstract—In this paper, we present a local, robot-centric navigation map optimized for autonomous mobile robots operating in unknown environments, enhancing their onboard perception systems for collision-free operation with far look-ahead distances. Utilizing a novel converging covariance cell representation, our approach effectively analyzes hazards such as obstacles and hazardous slopes in both terrestrial and aerial navigation contexts. The new technique specifically targets mapping from stereo scenarios with ultra short baseline and highly oblique viewpoints close to the ground.

Our methodology surpasses traditional window-based hazard analysis by resolving sub-cell size obstacles and terrain gradients at the individual cell level, thereby avoiding the computational overhead typically associated with such analyses. It leverages a multi-resolution strategy adaptive to the range errors common in stereo vision systems, making it particularly suitable for embedded systems with computational limitations.

Functionality includes constant-time queries for height, obstacle presence, and slope details, boasting improvements in run time, memory usage, precision, and resolvable obstacle size compared to existing grid-based mapping algorithms. We validate our approach through rigorous simulation and real-world testing. This technique will be used for the local mapping and collision avoidance on NASA’s CADRE lunar rovers.

I. INTRODUCTION

Terrain perception and local mapping are crucial for autonomous mobile robots, ensuring safe, collision-free navigation. Specifically, autonomous off-road vehicles and planetary rovers need effective traversability analysis methods given the constraint of sensor accuracy, noise and limited available computational resources. Mainly surface obstacles like rocks, holes, or slopes are encountered, making a 2.5D map representation adequate. To identify hazards, 2.5D geometric grid maps assess discontinuities and slopes across multiple cells, marking obstacles on an occupancy grid map for navigation planning [1]. Within these maps, cell resolution is dependent on the sensor noise and limits the smallest detectable obstacle. High resolution maps become unreliable under high sensor noise while large cells fail to accurately represent obstacles.

The CADRE mission [2], employs short baseline stereo with an oblique, ground-near angle, generating a point cloud of uneven point density and measurement noise. High-density, low-noise data near the robot are scrutinized for small obstacles, while less precise data from further away identify larger obstacles for long-term planning. A multi layered map with different sizes and resolutions is used



Fig. 1. CADRE rover with the CCM map. Image is generated by overlaying the super-sampled map (5 cm/cell map, sampled at 1 cm). The piece-wise linear map cells are visible as small inclined surfels with color indicating the variance normal to the slope and thus obstacle presence.

to overcome this issue and match the sensor resolution in the map architecture. Different mapping techniques for the individual layers make optimal use of the available information and computational resources.

We present a new technique, estimating the local inclination as well as the intra-cell roughness for each cell as part of the measurement update. This allows for obstacle identification without geometric filtering and the detection of obstacles which are smaller than the resolution of the map. Fig. 1 illustrates an example of reconstructed terrain for a CADRE rover.

II. RELATED WORK

Robotic maps vary greatly in their configurations, ranging from basic binary occupancy maps [1], [3], [4] to intricate 3D representations based on Oct Trees and Voxels [5], [6]. While 3D maps are highly expressive, their use in path planning and collision avoidance is computationally intensive, leading to a preference for 2.5D maps [7], which represent height as a function of position.

While surface approximation can be achieved through methods including meshing [8] and piece-wise (bi-)linear functions [9], these maps pose difficulties in incremental updating. Fusion of new measurements into DEMs can use the latest measurement or mean of all heights, with sophisticated techniques employing one-dimensional Kalman Filters to consider sensor uncertainty during height estimation for each cell [10]–[12]. The setup in [12] can adapt to dynamic

¹ Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA papers@lennart-werner.de, { pproenca, roland.brockers } @ jpl.nasa.gov

² Computer Science XVII – Robotics, Julius-Maximilians-Universität Würzburg, Germany andreas.nuechter@uni-wuerzburg.de

resolution changes, assisting UAVs in hazard detection.

Kalman-based updates for robot-centric maps with an estimated uncertainty deterioration, based on robot movement, have shown promise in reflecting pose errors accurately in map uncertainty [13]. However, they assume minimal intra-cell terrain variance, which can be unrealistic for large cells encapsulating substantial surface features.

The Optimal Mixture of Gaussians (OMG) update overcomes this challenge by introducing a variance that approximates the height distribution of samples [14]. OMG provides a mean cell height and a confidence band in which the true terrain lies. Our work builds upon this idea by proposing a three-dimensional covariance per cell approach inspired by [15], facilitating the creation of a 2.5D map inclusive of slope and roughness data.

We align with the strategies in [16], [17] in utilizing plane fitting to terrain patches within each cell to estimate slope and roughness, critical for state-of-the-art traversability analysis. However, we advance beyond these methods by accommodating multi-size, multi-resolution mapping and facilitating incremental updates without necessitating iterative processes unfit for small embedded systems.

Traversability analysis often involves geometric processing of the height map, incorporating either rover simulation over a single resolution DEM to identify obstacles [18] or a roughness filter examining height differences in neighboring cells to assign roughness values [12]. Despite their applicability, these strategies are computationally expensive and confined to single resolution maps, hindering the detection of smaller obstacles [19]. Our approach aims to overcome these limitations, enhancing the efficiency and effectiveness of local traversability mapping.

III. CONTRIBUTION

This work introduces a *converging covariance map* (CCM) based representation with an incremental update for mobile robotics. State-of-the-art terrain mapping approaches assume piecewise-constant maps and therefore lose vital information for traversability analysis due to discretization. We propose an efficient piecewise linear map that allows to resolve smaller obstacles with larger cells and do precise slope estimation without additional cost. Our algorithm is designed to be executed in real-time on a small embedded system. Slope and variance are estimated in each cell, which enables reliable detection of obstacles smaller than the resolution of the map. A forgetting mechanism keeps the map adaptable to a changing environment or drift in pose estimation. Updating the map with new measurements is performed in linear time with point cloud size, similar to existing piece-wise constant maps. Buffering of measurements is not necessary. Supporting multi-size, multi-resolution map setups for efficient stereo vision based mapping, a pooling operation is formulated additionally to the incremental measurement update to incorporate measurements consistently and efficiently into multiple overlapping map layers. The combination of a traditional piece-wise constant map for the highest resolution and the new CCM map to estimate slope

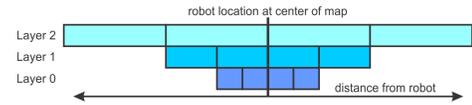


Fig. 2. Layout of multi-size, multi-resolution map. Areas in the center are covered by multiple resolutions.

and roughness for all low resolution layers forms a reliable and precise local map. The structure is shown in Fig. 2. Finally, the new map is evaluated and compared to state of the art techniques.

IV. CONVERGING COVARIANCE MAP

A. Motivation

Height measurement distribution within one cell is affected by both, the *Measurement Variance* of the sensor and the *Terrain Variance* due to irregularities of the surface. These variances combine to create the *Sample Variance* in each cell.

To enhance evaluations, it is beneficial to estimate the complete three-dimensional sample distribution, represented by $\Sigma_s \in \mathbb{R}^{3 \times 3}$, and the intra-cell mean of samples, denoted as $\mu_s \in \mathbb{R}^3$. This approach facilitates a detailed assessment compared to using traditional maps with constant pieces.

As illustrated in Fig. 3, utilizing the full covariance ellipsoid for each map cell aids in the accurate estimation of slope and surface-normal variance. This is essential for detecting obstacles effectively through the analysis of the ellipsoid's orientation and shape. Fig. 3 represents a section of the 3D covariance ellipse, showing one lateral axis alongside the elevation axis.

B. Cell Representation

In CCM, each cell requires five distinct elements from a three-dimensional covariance matrix, excluding the correlation between the two lateral axes which is not needed. The five (co)variance values describe variances across the three spatial axes and the covariances between the lateral and vertical axes. Moreover, to counteract unequal distribution of measurements and allow measurement weighting, each cell must maintain a record of a three-dimensional mean and one accumulated weight factor. Consequently, a total of nine values are stored for every CCM cell, namely the accumulated weight W , three-dimensional mean $\mu_{[x,y,z]}$ and five-dimensional (co)variance $\sigma_{[xx,yy,zz,xx,yz]}^2$. To simplify the incremental update, (co)variances are memorized as a W scaled value $S = W \cdot \sigma^2$, rather than being stored directly.

C. Incremental Update

1) *Update Equations*: To satisfy runtime and memory demands, an incremental update for (co)variances and means is essential. This update (eqs. (1) to (5)), facilitates the transition from the a-priori state at $t - 1$ to the posterior state t utilizing a single-point measurement. The respective cell must be identified through map discretization before initiating the update, using the measured coordinates denoted by $M_{[x,y,z]}$ and yielding the geometric cell center coordinates $C_{[x,y]}$. Calculated cell means are shown as $\mu_{c,[x,y,z]}$. The measurement weight, represented as w_t in (1), governs the

impact of the new data in relation to the existing dataset. This weight derives from the sensor model and stands as a crucial tuning parameter for the map, with a constant value proving sufficient in tested configurations.

Following the algorithm shown by [20], the incremental updates of mean and variances have been devised. This approach was extended to covariances in (5).

$$W_t = W_{t-1} + w_t \quad (1)$$

$$\mu_{m,t} = [M_x - C_x, M_y - C_y, M_z]^T \quad (2)$$

$$\begin{aligned} \mu_{c,t,[x,y,z]} &= \mu_{c,t-1,[x,y,z]} \\ &+ \frac{w_t}{W_t} \cdot (\mu_{m,t,[x,y,z]} - \mu_{c,t-1,[x,y,z]}) \end{aligned} \quad (3)$$

$$\begin{aligned} S_{[xx,yy,zz],t} &= S_{[xx,yy,zz],t-1} + w_t \cdot (\mu_{m,t,[x,y,z]} \\ &- \mu_{c,t-1,[x,y,z]}) \cdot (\mu_{m,t,[x,y,z]} - \mu_{c,t,[x,y,z]}) \end{aligned} \quad (4)$$

$$\begin{aligned} S_{[xz,yz],t} &= S_{[xz,yz],t-1} + w_t \cdot (\mu_{m,t,[x,y]} \cdot \mu_{m,t,[z,z]} \\ &- \mu_{c,t-1,[x,y]} \cdot \mu_{c,t-1,[z,z]}) \\ &+ W_t \cdot (-\mu_{c,t,[x,y]} \cdot \mu_{c,t,[z,z]} \\ &+ \mu_{c,t-1,[x,y]} \cdot \mu_{c,t-1,[z,z]}) \end{aligned} \quad (5)$$

2) *Weight Limiting*: To ensure consistent significance for new measurements in the incremental update, a convergence limiter is implemented. This step is essential as the current setup tends to diminish the relevance of new data over time, with the accumulated weight, W , increasing through straightforward addition (1). This growth in W can also lead to float precision issues, causing S values to escalate to a point where they cannot be precisely represented.

This limiter activates when the accumulated weight exceeds a specified threshold, W_l . In this case, the procedures outlined in eqs. (6) to (7) are executed to maintain the relative weight of new measurements and keep the map's adaptability to evolving surroundings and resilience towards pose estimation drift. It essentially ensures a controlled response to environmental changes by dictating the influence of new observations after full convergence.

$$W_{t+1} = W_l \quad (6)$$

$$S_{[xx,yy,zz,xz,yz],t+1} = S_{[xx,yy,zz,xz,yz],t} \cdot W_l / W_t \quad (7)$$

D. Map Query

A CCM map query provides the cell's mean height, two-axis inclination, and surface normal variance, allowing for sub-cell resolution estimations of height and inclination. The procedure for querying the X and Y axis inclinations (a , b), along with the elevation (h_q) at the query point designated as $([q_x, q_y]^T)$, is presented in eqs. (8) to (10). The queries utilize the Euclidean center coordinates of the queried cell, denoted as $C_{[x,y]}$. Furthermore, the surface-normal variance σ_s^2 query is calculated by equation (11). The total surface inclination is expressed by $i_{total} = \sqrt{a^2 + b^2}$. These metrics are vital in assessing traversability through threshold comparison, a

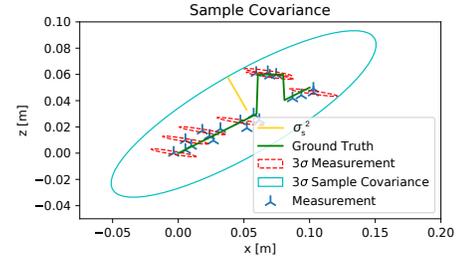


Fig. 3. 2D sliced CCM cell: Visualizing measurements, stereo measurement uncertainty and sample covariance. The cell contains an obstacle.

process detailed in section IV-E.

$$q_c = [q_x - C_x, q_y - C_y]^T \quad (8)$$

$$a = \frac{\sigma_{xz}^2}{\sigma_{xx}^2}, b = \frac{\sigma_{yz}^2}{\sigma_{yy}^2}, c = \mu_{c,z} - a \cdot \mu_{c,x} - b \cdot \mu_{c,y} \quad (9)$$

$$h_q = c + a \cdot q_{c,x} + b \cdot q_{c,y} \quad (10)$$

$$\sigma_s^2 = \sigma_{zz}^2 - \frac{\sigma_{xz}^4}{\sigma_{xx}^2} - \frac{\sigma_{yz}^4}{\sigma_{yy}^2} \quad (11)$$

Sub-cell elevation interpolation presumes an equal sample distribution within a cell, a condition disrupted by shadowing. As a result, partially covered cells are labeled as *unknown*. A cell of width w is considered adequately covered if the lateral standard deviations exceed a user-defined fraction, τ , of the maximum theoretical value as given in (12). In this study, setting $\tau = 0.8$ yielded good results.

$$\sigma_{xx}, \sigma_{yy} \geq \tau \cdot \sqrt{\frac{w^2}{12}} \quad (12)$$

E. Normal Variance Thresholding

The dependency between surface-normal variance and obstacle presence is illustrated in Fig. 3. The derived surface normal variance, denoted as σ_s^2 , portrays the *bulkiness* of the ellipsoid, aiding in sub-cell discontinuity detection. The variance threshold for a given obstacle height is determined through simulation.

A ground truth step of the minimum obstacle height is simulated, followed by computing the surface normal variance for increasing obstacle cell coverage levels. The resulting variance is exemplified in Fig. 4 for 32 cm cells with obstacles, equivalent to 75% of the cell width (24 cm). Setting the threshold to the variance of 50% cell coverage enables the detection of obstacles covering 11% to 91% of the cell, although the threshold can be modified to suit different scenarios, sensor configurations or expected obstacles.

F. Pooling

Pooling serves as a crucial strategy in significantly reducing the computational burden in multi-layer, multi-size mapping algorithms. It allows new measurements to be incorporated solely at the highest available resolution, bypassing the necessity of updating all layers; the lower resolution layers are only updated at the time of query through the fusion of high-resolution cells. However, given that individual CCM cells estimate distinct distributions characterized by unique

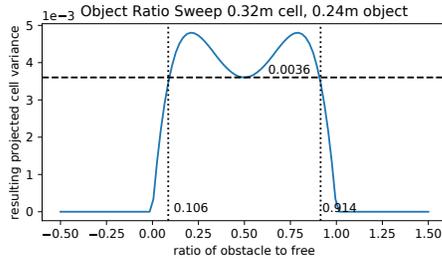


Fig. 4. Cell variance for obstacle coverage levels.

means and covariances, merging them into a single entity using traditional means is not feasible. Standard Gaussian mixture approaches, which operate under the assumption of all measurements observing a uniform value, are inapplicable here [14]. The appropriate strategy to pool four higher-resolution CCM cells (references with subscript c) into a single lower-resolution cell (subscript C) is outlined in eqs. (13) to (15) and is based on the combination of partially observed distributions. The additional index i identifies one of the small cells which are to be merged.

$$\mu_C = \frac{\sum_{i=0}^4 \mu_{c,i}}{4} \quad (13)$$

$$\sigma_{C,[x,y,z]}^2 = \frac{1}{4} \cdot \left[\sum_{i=0}^4 \sigma_{c,[x,y,z],i}^2 + (\mu_{C,[x,y,z]} - \mu_{c,i,[x,y,z]})^2 \right] \quad (14)$$

$$\sigma_{C,[xz,yz]}^2 = \frac{1}{4} \cdot \left[\left(\sum_{i=0}^4 \sigma_{c,[xz,yz],i}^2 + (\mu_{C,[x,y]} - \mu_{c,i,[x,y]}) \cdot (\mu_{C,[z]} - \mu_{c,i,[z]}) \right) \right] \quad (15)$$

G. Size Considerations

Since CCM estimates the measurement sample distribution, it is vital to pay attention to the stereo sensors' measurement noise. The chosen cell size must be sufficiently large to prevent range noise from dominating the variance induced by terrain features within a cell. Otherwise, CCM will estimate values influenced more by the camera's incidence angle rather than the ground truth inclination and roughness.

To ensure the dominant axes in the covariance reflect the surface rather than sensor noise, one must compare the measurement standard deviation (σ_m) with the anticipated standard deviation of a uniform distribution of cell dimensions, defined as $\sigma_r = \sqrt{\frac{w^2}{12}}$, where w denotes the cell width. A CCM convergence to the terrain values requires σ_r to be at least twice as large as σ_m for reliable rejection of the range noise. Consequently, this constrains the cell size with a lower bound for a given distance and sensor model.

V. EXPERIMENTAL EVALUATION

We evaluate the CCM algorithm's performance relative to existing piece-wise constant mapping techniques, verifying its advantages in precision and sub-cell obstacle detection. We chose the in elevation mapping most commonly used Kalman update and the more advanced and sample-variance aware OMG update as representatives for the piece-wise constant maps to benchmark against. We assess the algorithms based on their terrain representation accuracy and

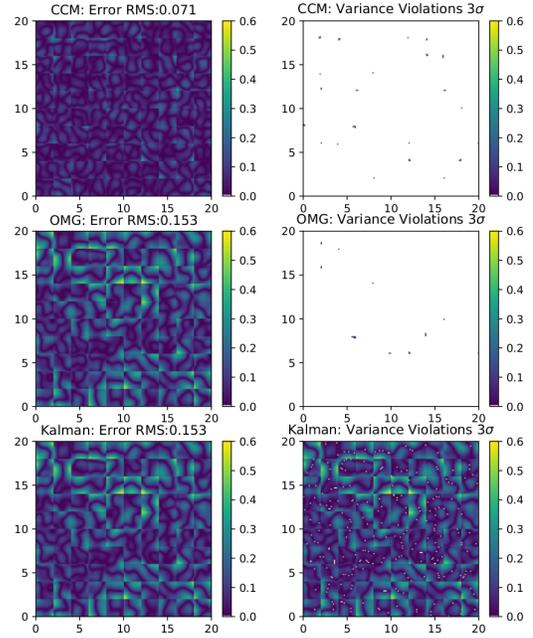


Fig. 5. 20 x 20m map. CCM, OMG, KF at 3 m resolution. Evaluation of precision (left) and variance representation (right). Precision: color indicates error between real terrain and map estimate. Variance: White indicates areas where elevation error $< 3\sigma$. Color quantifies error outside these bounds.

hazard detection capabilities. Here, hazards refer to areas exceeding a specified local inclination or featuring terrain discontinuities beyond the acceptable limit.

A. Reconstruction Precision and Variance

This section evaluates precision and confidence using synthetic data from a noise-added point cloud based on simulated terrain, as shown in Fig. 5.

In a 3 m resolution test, there's a clear distinction in elevation error plots on the left, showing the absolute difference between actual and estimated elevations and revealing finer details beneath the cell level. This difference underscores the discrepancy between piece-wise linear and constant mapping techniques. Both Kalman and OMG updates produce identical elevation error maps and total RMS errors due to their identical elevation estimation processes. The CCM map, however, displays significantly lower errors in individual cells, leading to a lower overall elevation reconstruction error (RMS error of 0.07 vs. 0.15 in piece-wise constant maps). Further tests indicate that Piece wise Constant maps and CCM with half resolution achieve similar continuous terrain precision, maintaining comparable memory usage.

The right column's variance violation plots assess the variance accuracy, highlighting areas where the actual surface falls outside the estimated variance's 3σ band. The Kalman update fails to accurately represent sample variance, while OMG and CCM accurately predict sample variance. However, OMG's utility is limited due to its sensitivity to both inclination and surface roughness when estimating elevation variance. In contrast, CCM uniquely estimates variance orthogonal to the first-order fit, separating slope from roughness and resulting in lower variance figures.

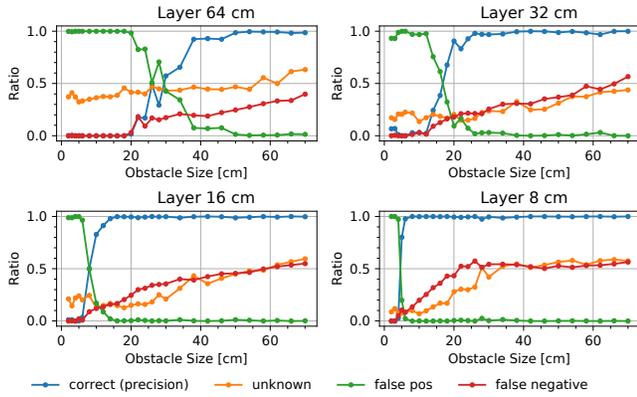


Fig. 6. Obstacle detection performance CCM

B. Obstacle Detection

The CCM algorithm’s obstacle detection capability was evaluated through high-fidelity simulations and real-world data, focusing on scenarios from NASA’s CADRE mission. CADRE rovers use ultra-short baseline stereo cameras with a 5 cm separation, positioned 17.5 cm above the ground. The CCM layers are calibrated for obstacles sized at 75% of a cell, as outlined in section IV-E. For comparison, roughness filter thresholds for contemporary maps activate at a 20 deg slope when set to the lowest effective value.

A series of Monte-Carlo simulations assessed map performance across various lunar terrains, paths, and obstacle sizes using a block-matching stereo vision algorithm for precise depth and noise evaluation, closely replicating actual rover setups. In these simulations, cell classifications at each timestep were recorded as *correct*, *false positive* (missed obstacle), *false negative* (incorrectly marked obstacle, yet passable), or *unknown* (not observed), determined by their traversal characteristics and real obstacle presence.

As depicted in Fig. 6, each data point aggregates the time-summed areas of correctly and incorrectly classified regions for distinct obstacle dimensions. The graph shows a dip in false positives as obstacle size escalates, identifying the obstacle detection capacity across different resolutions.

A partial detection of an obstacle still causes a visible false positive rate which impacts the measured precision. The small residual false positive rate after the main drop is caused by stereo border effects and do not pose a drive hazard as explained in Section V-E. Due to the conservative tuning of safety margins, a false negative of only 50% is achieved. The false negative areas are all adjacent to a real obstacle.

Obstacle detection in piece-wise constant maps is constrained by the need for the roughness filter. As shown in Fig. 7 and supported by [19], reliable detection occurs when the obstacle size is approximately $3\times$ the map resolution. This ensures correct obstacle height representation in the map.

Evaluating slope in a piece-wise constant map can also be done using a roughness filter or by fitting a plane to multiple cells. Both techniques are computationally intensive tasks. The roughness filter threshold, influenced by slope, thus cannot be arbitrarily lowered to detect smaller obstacles without increasing false negative rates in inclined areas. These

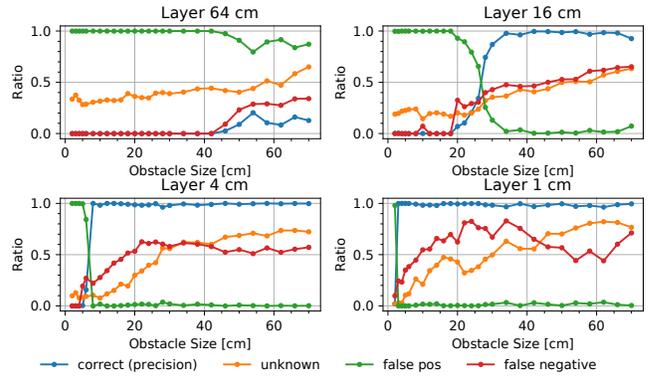


Fig. 7. Obstacle detection performance OMG.



Fig. 8. Picture of the test setup at JPL.

operations, supplementary to map updates and roughness estimation, are unnecessary with the CCM algorithm.

C. Real World Tests

Real-world tests were performed to validate simulation results, utilizing various sized rock and crater obstacles situated in the JPL Mini Marsyard, depicted in Fig. 8 and Fig. 9. An external optical pose tracking system ensured evaluation independence from localization, with seven Vicon Vantage cameras [21] delivering pose updates at 100 Hz. A survey-grade laser scanner served as the ground truth for elevation and obstacle assessments. Common targets in the Vicon data and laser scans are used for the alignment of ground truth elevation data and the odometry frame. This enables precise qualitative and quantitative evaluation of precision and obstacle detection performance.

The same evaluation pipeline applied to simulated data was used in real-world experiments to maintain consistency. Real-world findings match the simulated results, affirming the relative terrain reconstruction precision between the maps and the reproducibility of obstacle detection performance in real settings. The proposed algorithm and setup runs on the target hardware (Snapdragon 820) with $850\text{px} \times 650\text{px}$ stereo at 80 ms update time and a refresh rate of 5 Hz.

Fig. 9 shows a terrain patch mapped by CCM and OMG. The top row compares the DEM smoothness at an identical map resolution, while the bottom row highlights the accurate detection of a small obstacle by an 8 cm CCM layer — a detail not captured by the piece-wise constant map with double the resolution.

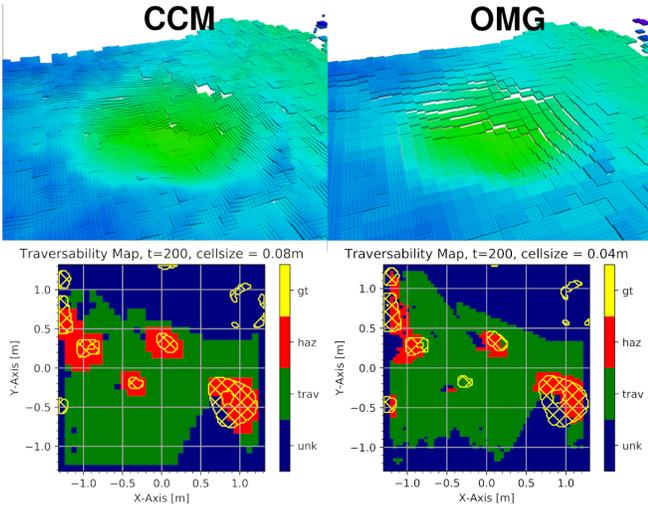


Fig. 9. Real world data CCM (left column) and OMG (right column). Top row: Terrain reconstruction of a crater (both maps in top row resolution: 8 cm for visual comparison). Bottom row: OMG with twice the resolution (4 cm) does not detect an obstacle CCM finds at 8 cm.

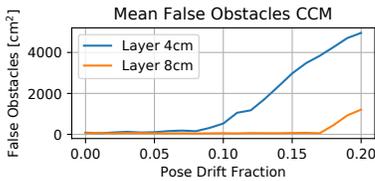


Fig. 10. Pose Drift Resilience CCM.

D. Drift Resilience

The integration cap W_l greatly affects drift resilience. In piece-wise constant maps, obstacles can drift outside labeled areas, a problem compensated for by averaging. Conversely, CCM mistakenly marks safe terrain as hazardous due to heightened sample variance arising from pose estimate drift.

We assessed drift resilience via simulation on wave shaped terrain, to represent the worst case scenario. Fig. 10 presents the findings for 4 cm and 8 cm resolution maps. The chosen W_l withstands drift by 8% for 4 cm and 16% for 8 cm, proving robustness to anticipated pose estimation drift.

E. Limitations

Depth errors from stereo vision predominantly escalate at object borders, leading to point cloud streaks. This is especially fatal at the top rim of an obstacle, where incorrect depth values extend measurements far behind the real obstacle. Since areas behind an obstacle lack measurements, these false data points significantly affect those cells, forming a flat surface in the obstacle’s shadow. This can potentially misclassify a cell as traversable if it surpasses the minimal measurement threshold. However, since this phenomenon is confined to an obstacle’s rim and extends behind it, it does not pose an immediate driving risk.

VI. APPLICATIONS

Low-resolution maps benefit from the discussed technique, as high measurement noise and sparse sensor data render

high-resolution maps unsuitable. Traditional geometric filtering proves inadequate for obstacle detection here.

Developed primarily for NASA’s CADRE mission, the CCM algorithm within a multi-size, multi-resolution mapping framework, reliably detects obstacles over 48 cm at a 5 m detection range. The strategy suits the mission’s unique setup involving a small stereo camera baseline and a ground clearance, which causes variable error variance in depth images. These factors, coupled with restrictions on map size, minimum obstacle detectability, and slope traversability, negate the utility of conventional piecewise constant maps.

The presented map can be used with different types of 3D sensors such as Lidar, Radar or SFM sensors. The sensors measurement noise model is required to find the correct cell resolution for the desired map size.

VII. CONCLUSION

CCM serves as an effective tool for simultaneously estimating the height and slope of a terrain, presenting a viable solution for computationally limited mobile robotic platforms that require both slope and hazard data for secure operations. It outperforms traditional methods utilizing piecewise constant elevation maps in precision, while consuming comparable computational resources.

The multi-resolution methodology aptly addresses the increasing range uncertainty encountered in 3D reconstructions generated by stereo camera systems, finding substantial applicability in autonomous vehicles. Presently, we are adapting this approach for NASA’s upcoming CADRE mission scheduled for lunar deployment in 2024 [22].

VIII. FUTURE WORK

Currently, CCM integrates pre-aligned point clouds from stereo vision into a unified map representation, utilizing a sample covariance distribution in each cell to generate probability predictions for incoming measurements. This framework could evolve to incorporate pose estimate corrections based on the derived observation probabilities, mitigating pose drift. Furthermore, the covariance parameters can foster quicker map patch matches, enabling global mapping and patch stitching through efficient covariance cell alignments, a strategy that promises to be faster than the conventional iterative closest point (ICP) algorithm used in map alignments.

Full 3D reconstruction with the presented technique would also be feasible, but requires a different voxel based memory setup and slight modification of the measurement update.

Due to the high obstacle detection performance and low computational load, JPL will investigate the usability of the presented map for landing site detection in future martian areal vehicles.

ACKNOWLEDGMENT

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the JPL Visiting Student Research Program (JVSRP) and the National Aeronautics and Space Administration (80NM0018D0004).

REFERENCES

- [1] A. Elfes, "Occupancy grids: A probabilistic framework for robot perception and navigation," *Tese de doutorado, Electrical and Computer Engineering, Carnegie Mellon University*, 1989.
- [2] [Online]. Available: <https://www.jpl.nasa.gov/missions/cadre>
- [3] J. Borenstein and Y. Koren, "Real-time map building for fast mobile robot obstacle avoidance," in *Mobile Robots V*, Mar 1991, p. 74–81.
- [4] D. Joubert, W. Brink, and B. Herbst, "Pose uncertainty in occupancy grids through monte carlo integration," in *IEEE ICRA*, Nov 2013, p. 1–6.
- [5] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: an efficient probabilistic 3d mapping framework based on octrees," *Autonomous Robots*, vol. 34, p. 189–206, Apr 2013.
- [6] H. Oleynikova, Z. Taylor, M. Fehr, J. Nieto, and R. Siegwart, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," *IEEE IROS*, p. 1366–1373, Sep 2017.
- [7] M. Herbert, C. Caillas, E. Krotkov, I. Kweon, and T. Kanade, "Terrain mapping for a roving planetary explorer," in *IEEE ICRA*, 1989, p. 997–1002.
- [8] J. Zienkiewicz, A. Tsiotsios, A. Davison, and S. Leutenegger, "Monocular, real-time surface reconstruction using dynamic level of detail," in *3DV*, Oct 2016, p. 37–46.
- [9] J. Baltes and J. Anderson, "Interpolation methods for global vision systems," in *RoboCup 2004: Robot Soccer World Cup VIII*. Springer, 2005, p. 434–442.
- [10] L. B. Cremean and R. M. Murray, "Uncertainty-based sensor fusion of range data for real-time digital elevation mapping (rtDEM)," *IEEE ICRA*.
- [11] C. Forster, M. Faessler, F. Fontana, M. Werlberger, and D. Scaramuzza, "Continuous on-board monocular-vision-based elevation mapping applied to autonomous landing of micro aerial vehicles," in *IEEE ICRA*, May 2015, p. 111–118.
- [12] P. Schoppmann, P. F. Proença, J. Delaune, M. Pantic, T. Hinzmann, L. Matthies, R. Siegwart, and R. Brockers, "Multi-resolution elevation mapping and safe landing site detection with applications to planetary rotorcraft," *IEEE IROS*, Nov 2021.
- [13] P. Fankhauser, M. Bloesch, C. Gehring, M. Hutter, and R. Siegwart, "Robot-centric elevation mapping with uncertainty estimates," in *CLAWAR*, Sep 2014.
- [14] P. F. Proença, J. Delaune, and R. Brockers, "Optimizing terrain mapping and landing site detection for autonomous uavs," *IEEE ICRA*, p. 7, 2022.
- [15] C. Park, S. Kim, P. Moghadam, C. Fookes, and S. Sridharan, "Probabilistic surfel fusion for dense lidar mapping," *ICCV*, Sep 2017.
- [16] S. Singh, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, and K. Schwehr, "Recent progress in local and global traversability for planetary rovers," in *IEEE ICRA*, Apr 2000, p. 1194–1200 vol.2.
- [17] D. B. Gennery, "Traversability analysis and path planning for a planetary rover," *Autonomous Robots*, no. 6, p. 131–146, 1999.
- [18] M. Daily, J. Harris, D. Keirse, D. Olin, D. Payton, K. Reiser, J. Rosenblatt, D. Tseng, and V. Wong, "Autonomous cross-country navigation with the alv," in *IEEE IRA*, Apr 1988, p. 718–726.
- [19] R. Brockers, J. Delaune, P. Proença, P. Schoppmann, M. Domnik, G. Kubiak, and T. Tzanetos, "Autonomous safe landing site detection for a future mars science helicopter," in *IEEE Aerospace Conference*, Mar 2021, p. 1–8.
- [20] T. Finch, "Incremental calculation of weighted mean and variance," *University of Cambridge Computing Service*, p. 8, 2009.
- [21] [Online]. Available: <https://www.vicon.com/>
- [22] [Online]. Available: https://www.nasa.gov/directorates/spacetech/game_changing_development/projects/CADRE