

Lunar ROADSTER: Autonomous Bulldozer for Lunar Terrain Manipulation

Bhaswanth Ayapilla^{*‡}, Boxiang Fu^{*‡}, Ankit Aggarwal^{*}, Deepam Ameria^{*}, Simson D’Souza^{*}, John M. Dolan, and William “Red” Whittaker

I. INTRODUCTION

Humanity is preparing to return to the Moon, with particular interest in the lunar South Pole [1] [2], a region containing water ice that could support human life. Roadwork and terrain shaping for infrastructure may come to the Moon over time, but, as on Earth, the earliest manifestations will likely improve rocky, cratered terrain into traversable trails. One compelling use case is sun-synchronous circumnavigation — continuously outpacing the lunar sunset to maintain solar power — a feat feasible at speeds as low as 2.5 kph at 81° latitude (Table I). Such a mission demands continuously traversable terrain along its path, motivating autonomous trail grooming: filling craters by pushing rim regolith into their interiors to reduce surface slope.

TABLE I: Average Speed Required for Sun-Synchronous Circumnavigation at High Lunar Latitudes

Latitude	Distance (km)	Speed (kph)
70°	3,700	6
75°	2,800	4
80°	1,870	3
81°	1,529	2.5

This paper presents Lunar ROADSTER, a terrestrial prototype of an autonomous bulldozer for lunar trail grooming with minimal human supervision (Figure 1). Unlike flat-grading tasks addressed in prior work [3], trail grooming along a circumnavigation route requires a system that can selectively identify, navigate to, and manipulate craters along a reference latitude, all while traversing a closed circular path around the worksite.

Our primary contributions are: (1) a GPS-free localization architecture fusing a fisheye ceiling-light tracker with IMU and wheel encoders, with the ceiling-light tracking modality directly analogous to stellar navigation on the Moon, (2) a real-time, ground-level crater perception pipeline using YOLOv8-nano for metric geometry extraction and manipulation pose generation, and (3) a closed-loop demonstration of autonomous crater detection, navigation, grooming, and revalidation across multiple craters.

II. METHODOLOGY

A. System Overview

The system receives a global 2D map of the worksite and classifies craters as gradable targets or obstacles to avoid.

^{*}Equal Contribution. [‡]Corresponding Authors.

All authors are with Carnegie Mellon University, Pittsburgh, PA, USA. {bayapill, boxiangf}@cs.cmu.edu



Fig. 1: Lunar ROADSTER Rover Platform

It navigates to each gradable crater, detects its geometry onboard using a custom YOLOv8-nano perception pipeline [4], and executes a frontal dozing and backblading sequence to push regolith from the crater’s rim, into its interior. After each pass, a depth-based validation module evaluates terrain slope. If the groomed surface exceeds the threshold, the system replans and executes another pass. This closed-loop grooming-validation cycle continues until the trail segment is confirmed traversable, at which point the rover advances to the next crater. The system is validated against quantitative traversability and ground pressure requirements derived from lunar regolith simulant constraints.

B. Hardware

The bulldozer is designed to operate safely on loose regolith and manipulate terrain. The system is designed for a contact pressure of 1.37 kPa to minimize sinkage, with a total mass of 23.8 kg. The aluminum extrusion chassis features a 0.5 m wheelbase and a 0.43 m track width, along with a roll-averaging rocker suspension that generates 138 Newtons of drawbar pull on flat terrain. For locomotion, four planetary-gear DC motors drive the front and rear differentials, while a double Ackermann steering configuration provides independent control of both axles. The platform uses 0.22 m diameter 3D-printed PLA wheels, optimized to maximize drawbar pull, providing sufficient traction for both nominal traversal and grooming loads.



Fig. 2: BEN State Transition Flow

The dozer assembly comprises a front-mounted blade, aluminum support arms, and a feedback-controlled linear actuator. Machined from Al 6061, the blade measures 600 mm by 152 mm and provides a 0.09-square-meter swept area to increase material throughput per pass. Blade elevation is managed by an Actuonix L16-P actuator with 200 Newtons of lift, yielding a 20.6 lift-to-weight safety factor over the assembly weight. A potentiometer provides feedback on the blade’s angular position, enabling closed-loop control.

The platform has several sensors: an upward-facing fish-eye camera (SkyCam) mounted on a three-axis mechanical gimbal, a VN100 IMU, a ZED 2i stereo camera, and motor encoders. A Jetson Orin runs the full ROS2 autonomy stack, fusing SkyCam, ZED 2i, IMU data, and wheel encoders to drive navigation, perception, and manipulation decisions. An Arduino Due is used for lower-level hardware control and bridges the ROS2 stack with the drive hardware. The Arduino commands the dozer blade elevation via an Actuonix control board and controls the drive and steering through a Basicmicro RoboClaw controller with quadrature encoder feedback.

Three 20 V batteries power the system via a distribution board with a 20-12 V DC-DC buck converter, with a fuse, wireless power switch, and physical emergency stop ensuring system safety.

C. Software

The autonomy stack runs on ROS2 Humble on the Jetson Orin. A Behavior Executive Node (BEN) (Figure 2) implements a finite state machine that sequences the mission through global navigation, crater perception, manipulation planning, grooming execution, and terrain revalidation. A centralized multiplexer node supports four control modes: IDLE, FULL_TELEOP, AUTOGRADER, and FULL_AUTONOMY, enabling clean handoff between operator and autonomy stack without restarting any nodes.

On the Moon, orbital elevation maps (e.g., LRO LOLA [5]) provide the terrain topography prior needed for mission planning. In our testbed, a FARO 3D laser scan serves the equivalent role: the point cloud is height-thresholded offline to produce a 2D occupancy grid marking crater rims as occupied and traversable floor as free. This map remains static throughout execution.

The global planner first classifies craters by diameter: those below a threshold d_{thresh} are designated gradable targets, while larger craters are marked as obstacles in the occupancy grid. Since the mission targets gradable craters along a designated lunar latitude, their centroids are fitted to a circle using linear least-squares, refined by Gauss-Newton iteration with Huber weighting to reduce sensitivity to outlier craters. The sorted centroids are connected by a closed Catmull-Rom spline [6] providing C^1 -continuous arc-length parameterization, from which arbitrary poses can be projected to obtain arc coordinate s , tangent heading $\Psi_{\text{ring}}(s)$, and signed lateral offset d_{\perp} . A lattice-based A* search [7] in $SE(2)$ plans paths along this ring using straight, left, and right motion primitives, each advancing the rover by a fixed step length Δs along the resulting arc. States are discretized into a (x, y, θ) grid with heading binned into 16 sectors, and the cost function penalizes path length, steering effort, lateral ring deviation, heading misalignment, forward progress along the ring, and obstacle proximity. Sub-goals placed at uniform arc-length intervals between consecutive crater pairs keep planning tractable. For short-range grooming motions, a lightweight 2D-grid A* over the same occupancy map with an 8-connected neighborhood generates forward dozing and backblade trajectories. Both regimes are tracked by a single Pure Pursuit controller [8], parameterized with a tighter look-ahead distance and goal tolerance for manipulation mode.

The perception subsystem detects craters in real-time using a custom YOLOv8-nano model trained on manually annotated RGB images of the lunar testbed (Figure. 3). Deployed with CUDA acceleration on the Jetson, the model achieves detection confidence between 55–95% on the live ZED 2i feed, driven by variation in crater size, viewing angle, and lighting. RGB bounding boxes of the craters are fused with registered depth data to extract each crater’s centroid position (x, y, z) and diameter d in the robot frame, with measurements smoothed over 15 successive frames at 1 Hz to mitigate the effect of low-confidence detections on geometry extraction. These are used to compute grooming poses (source, sink, and backblade) placed at opposite ends of the crater along the rover’s approach heading. Blade height is dynamically computed from depth data at the corresponding ground location, ensuring only the necessary regolith is displaced per pass and preventing excessive accumulation or wheel slip. The validation subsystem, invoked after each grooming pass, aggregates ZED 2i point clouds into voxels, estimates per-voxel surface normals via PCA over k -nearest neighbors [9], and infers the slope. If the slope exceeds 5° , the BEN re-invokes perception and re-plans the grooming sequence.

Localization fuses fisheye ceiling-light tracking (SkyCam) with IMU and wheel encoder data via an Extended Kalman Filter (EKF) [10]. The SkyCam is mounted on a three-axis mechanical gimbal pointed vertically upward, which decouples camera orientation from rover roll and pitch to keep the ceiling-light grid undistorted during traversal over uneven terrain. The fisheye camera’s contrast and saturation is custom-tuned to allow feature detection in both bright

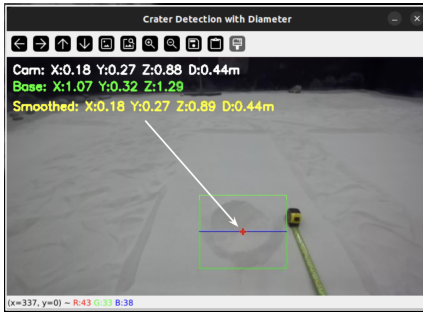


Fig. 3: Crater Detection and Geometry Extraction



(a) Before (slope = 14.78°) (b) After (slope = 1.44°)

Fig. 4: Terrain Before (Left) and After (Right) a Single Autonomous Crater Grading Run.

and dark environments. Detected ceiling light features are projected into a normalized frame using the calibrated fisheye inverse projection model. Their positions are then used to solve a nonlinear least-squares problem via Levenberg-Marquardt [11] over a periodic ceiling grid, recovering planar translation (u, v) , and rotation θ relative to the grid. These are converted to metric world coordinates by scaling with the known ceiling height h : $X = -u \cdot h, Y = -v \cdot h, \psi = -\theta$, and fused with IMU orientation in the EKF to produce the full rover pose.

III. EXPERIMENTAL RESULTS

All experiments were conducted in a $6\text{m} \times 6\text{m}$ indoor testbed of a simulated lunar surface environment with distributed craters of varying sizes. Each trial consisted of a full autonomous mission: grading ring construction, path planning, crater detection, and execution of the grooming sequence, including forward dozing and backblading passes. An example illustrating the terrain before and after a single crater grading run is shown in Figure 4.

System performance was evaluated against three mission performance requirements established at the outset of the project:

- **M.P.1:** Planned path deviation from reference latitude $\leq 25\%$
- **M.P.2:** Actual path-following deviation $\leq 10\%$
- **M.P.3:** Wheel contact pressure $\leq 1.5\text{kPa}$

A. Planned Path Deviation (M.P.1)

Path deviation measures how closely the A*-planned path follows the ideal crater ring, computed as the mean radial deviation of the planned path from the best-fit circle through all crater centroids, normalized by the reference latitude

length of 14.66 m (radius 2.33 m). Table II shows the mean deviation across all trials (22.89%) satisfies the $\leq 25\%$ requirement in all five trials.

B. Path-Following Deviation (M.P.2)

Path-following deviation measures how closely the rover's actual traversed trajectory tracked the planned path, computed as the normalized difference between planned and actual traversed path lengths. Table III shows the mean deviation across five trials (7.01%) satisfies the $\leq 10\%$ requirement.

C. Contact Pressure (M.P.3)

The Bekker pressure-sinkage model [12] is used to calculate contact pressure. Assuming a 23.8 kg vehicle mass distributed across four wheels in 1.62m/s^2 lunar gravity, the per-wheel load is 9.72 Newtons. The environment is modeled using Apollo Type B regolith parameters: a cohesive modulus of 1.4×10^3 , a frictional modulus of 8.2×10^5 , and a deformation exponent of 1.2 [13]. Given wheel dimensions of 0.11 m in radius and width, solving the equilibrium equation for wheel load versus sinkage yields the contact pressure as 1.37 kPa, below the 1.5 kPa threshold.

IV. CONCLUSION

This paper presented an autonomous bulldozer system for grooming traversable surface trails to enable sun-synchronous circumnavigation on the Moon. We demonstrated a closed-loop pipeline integrating real-time crater detection via a custom YOLOv8-nano model, kinematically-constrained path planning, regolith manipulation, and terrain revalidation in a simulated lunar testbed. The system met all three mission performance requirements: mean planned path deviation of 22.89% ($\leq 25\%$), mean path-following deviation of 7.01% ($\leq 10\%$), and wheel contact pressure of 1.37 kPa ($< 1.5\text{kPa}$). A SkyCam-based localization architecture, fusing fisheye ceiling-light tracking with IMU and wheel odometry via EKF provided continuous pose estimation. On the

TABLE II: Planned Path Deviation from Crater Ring (M.P.1)

Trial	Planned Path (m)	Deviation (m)	Deviation (%)
1	11.09	3.57	24.34
2	11.51	3.15	21.47
3	11.11	3.55	24.20
4	11.51	3.15	21.47
5	11.29	3.37	22.97
Mean	11.30	3.36	22.89
Std. Dev.	0.20	0.21	1.34

TABLE III: Actual Path-Following Deviation (M.P.2)

Trial	Planned (m)	Actual (m)	Dev. (m)	Dev. (%)
1	11.09	9.65	1.44	12.98
2	11.51	12.97	1.46	12.68
3	11.11	11.42	0.31	2.79
4	11.51	10.76	0.75	6.52
5	11.29	11.28	0.01	0.09
Mean	11.30	11.22	0.79	7.01
Std. Dev.	0.20	1.18	0.59	5.22

lunar surface, this approach transfers naturally to star-tracker-based localization using the same algorithmic framework. Future work will address multi-DOF tool articulation for multi-directional crater approach and improved localization robustness.

ACKNOWLEDGMENT

We would like to express our gratitude to the individuals and groups whose guidance and support made the Lunar ROADSTER project possible. A special thanks to our Course Advisor, Dr. Dimitrios Apostolopoulos, for his academic oversight and feedback. The fabrication and assembly of the rover were greatly facilitated by Timothy Angert, our Workshop Supervisor. We extend our appreciation to Team CraterGrader and Team Lunar-X, our predecessor teams, for laying the foundational work that informed our system architecture. Additionally, we thank Dr. Wennie Tabib and Warren ‘Chuck’ Whittaker for their technical assistance, and Dr. David Wettergreen for providing access to the testing space required to validate our system on lunar-like terrain.

LIMITATIONS

While our system successfully demonstrated closed-loop autonomous crater grooming, there are a few limitations:

Localization: SkyCam-based localization provided continuous, GPS-denied pose estimation by tracking ceiling lights as proxies for stars, directly analogous to stellar navigation on the Moon. However, it remains sensitive to terrain-induced platform tilt that the mechanical gimbal does not fully compensate. The fisheye camera may also have limited effectiveness in intense sunlight and high albedo regolith environments that cause the sensor to become over-saturated. This motivates future work on software-based horizon correction and improved localization robustness.

Manipulation: The dozer assembly provides a single degree of freedom (vertical blade height), constraining the rover to approach each crater along a fixed heading. Craters that are misaligned with the rover’s approach direction cannot be groomed without additional global replanning. A blade that can turn left or right would allow the rover to approach from many more directions, increasing planning options and eliminating the need for the rover to line up.

Operational Overhead: Pre-mission setup requires manual mapping with a FARO laser scanner and external infrastructure deployment, currently taking several hours and limiting the frequency of field testing. However, lunar deployment would instead rely on orbital elevation maps such as LRO products to provide topography information.

Stuck-State Recovery: The BEN does not implement automated timeout or retry-limit handling for failed crater detections. If the system enters a stuck PERCEPTION state, operator intervention is required via a manual debug trigger to redirect the FSM. Automated stuck-state recovery is identified as a direction for future work.

REFERENCES

- [1] D. Folta and T. Sweetser, “Artemis mission overview: From concept to operations,” vol. 142, pp. 1629–1645, 01 2012.
- [2] Indian Space Research Organisation, “Chandrayaan-3 Details,” https://www.isro.gov.in/Chandrayaan3_Details.html, 2023, accessed: 2026-04-09.
- [3] R. Lee, B. Younes, A. Pletta, J. Harrington, R. Q. Wong, and W. R. Whittaker, “CraterGrader: Autonomous robotic terrain manipulation for lunar site preparation and earthmoving,” *arXiv preprint arXiv:2311.01697*, 2024. [Online]. Available: <https://arxiv.org/abs/2311.01697>
- [4] G. Jocher, A. Chaurasia, A. Stoken, and J. Liu, “Ultralytics yolov8,” <https://github.com/ultralytics/ultralytics>, 2023.
- [5] D. E. Smith, M. T. Zuber, G. B. Jackson, J. F. Cavanaugh, G. A. Neumann, H. Riris, X. Sun, R. S. Zellar, C. Coltharp, J. Connelly *et al.*, “The lunar orbiter laser altimeter investigation on the lunar reconnaissance orbiter mission,” *Space science reviews*, vol. 150, no. 1, pp. 209–241, 2010.
- [6] E. Catmull and R. Rom, “A class of local interpolating splines,” in *Computer aided geometric design*. Elsevier, 1974, pp. 317–326.
- [7] M. Pivtoraiko, R. A. Knepper, and A. Kelly, “Differentially constrained mobile robot motion planning in state lattices,” *Journal of Field Robotics*, vol. 26, no. 3, pp. 308–333, 2009. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.20285>
- [8] R. C. Coulter, “Implementation of the pure pursuit path tracking algorithm,” Tech. Rep., 1992.
- [9] R. B. Rusu and S. Cousins, “3D is here: Point Cloud Library (PCL),” in *IEEE International Conference on Robotics and Automation (ICRA)*. Shanghai, China: IEEE, May 9-13 2011.
- [10] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte, “A new approach for filtering nonlinear systems,” in *Proceedings of 1995 American Control Conference-ACC’95*, vol. 2. IEEE, 1995, pp. 1628–1632.
- [11] J. J. Moré, “The levenberg-marquardt algorithm: Implementation and theory,” in *Numerical Analysis*, G. A. Watson, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 105–116.
- [12] M. G. Bekker, *Theory of Land Locomotion: The Mechanics of Vehicle Mobility*. Ann Arbor, MI: University of Michigan Press, 1956.
- [13] G. H. Heiken, D. T. Vaniman, and B. M. French, *Lunar Sourcebook: A User’s Guide to the Moon*. Cambridge, UK: Cambridge University Press, 1991.