

# Lunar ROADSTER

## Autonomous Bulldozer for Lunar Terrain Manipulation

Bhaswanth Ayapilla<sup>†\*</sup>, Boxiang Fu<sup>†\*</sup>, Ankit Aggarwal<sup>†\*</sup>, Deepam Ameria<sup>†\*</sup>, Simson D'Souza<sup>†\*</sup>, John M. Dolan<sup>†</sup>, William 'Red' Whittaker<sup>†</sup>




### Problem Statement & Mission/Use Case


Humanity is preparing to return to the **Moon**, with the Artemis missions focusing on exploring the South Pole - a region rich in sites of interest. Establishing a circumnavigating route around the lunar pole will serve as a critical "highway" connecting these sites and enabling key activities such as transportation, human settlement, and resource extraction.

A solar-powered rover capable of sun-synchronous circumnavigation could achieve perpetual operation by avoiding lunar sunsets. At high latitudes, this is feasible at low speeds (2-5 kmph). However, these assumptions rely on the terrain being **flat and traversable**, free from major topographical challenges. A mission to **manipulate the lunar regolith in the circumnavigating path to make it more traversable** for future missions is thus a clear step forward. A robotic system can be designed to conduct these operations efficiently for extended durations.

**The Lunar Robotic Operator for Autonomous Development of Surface Trails and Exploration Routes** (Lunar ROADSTER) is an autonomous moon-working rover, capable of finding exploration routes and grooming the lunar surface to develop traversable surface trails. These groomed trails will become the backbone for the colonization of the Moon by enabling transportation, logistics, and enterprise development.



### Hardware Design & Testing Setup

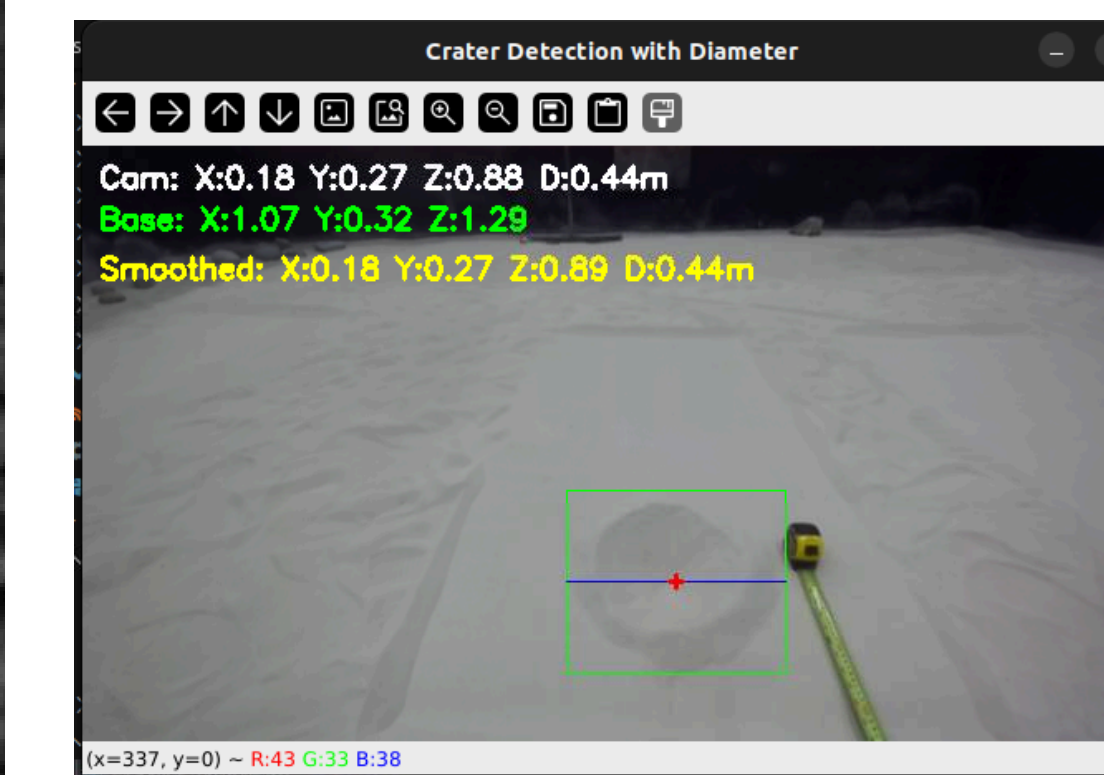


The ROADSTER platform is a **specialized, autonomous bulldozer** designed to execute high-force grading tasks in lunar-analogue environments. The mobility system comprises **135 kgf-cm drive actuators**, coupled with custom high-grip wheels engineered to maintain traction and maximize rimpull in loose regolith, while traversing uneven terrain. This drivetrain supports a **60 cm x 15 cm grading blade** and provides the necessary pushing power to displace significant soil volumes during leveling operations. To ensure operational endurance, the electromechanical architecture features reinforced tool actuation for consistent blade pressure and a streamlined power distribution system, which includes a custom PCB for reliable, high-current delivery.

**Test Site: Moon Yard, Planetary Robotics Laboratory, Carnegie Mellon University**

- Indoor Sandbox ~ 40m<sup>2</sup> area with Quikrete 1113 regolith simulant
- Task Space: 4 present craters and 1 spool crater
- 4.66m diameter circle represents the reference latitude
- Global Map obtained using a FARO Laser Scanner to emulate Lunar Reconnaissance Orbiter (LRO) Maps [1]

### Perception




The perception system is the rover's primary interface with the unstructured lunar environment, enabling it to autonomously identify work sites, adapt its tool control, and verify the quality of its construction tasks.

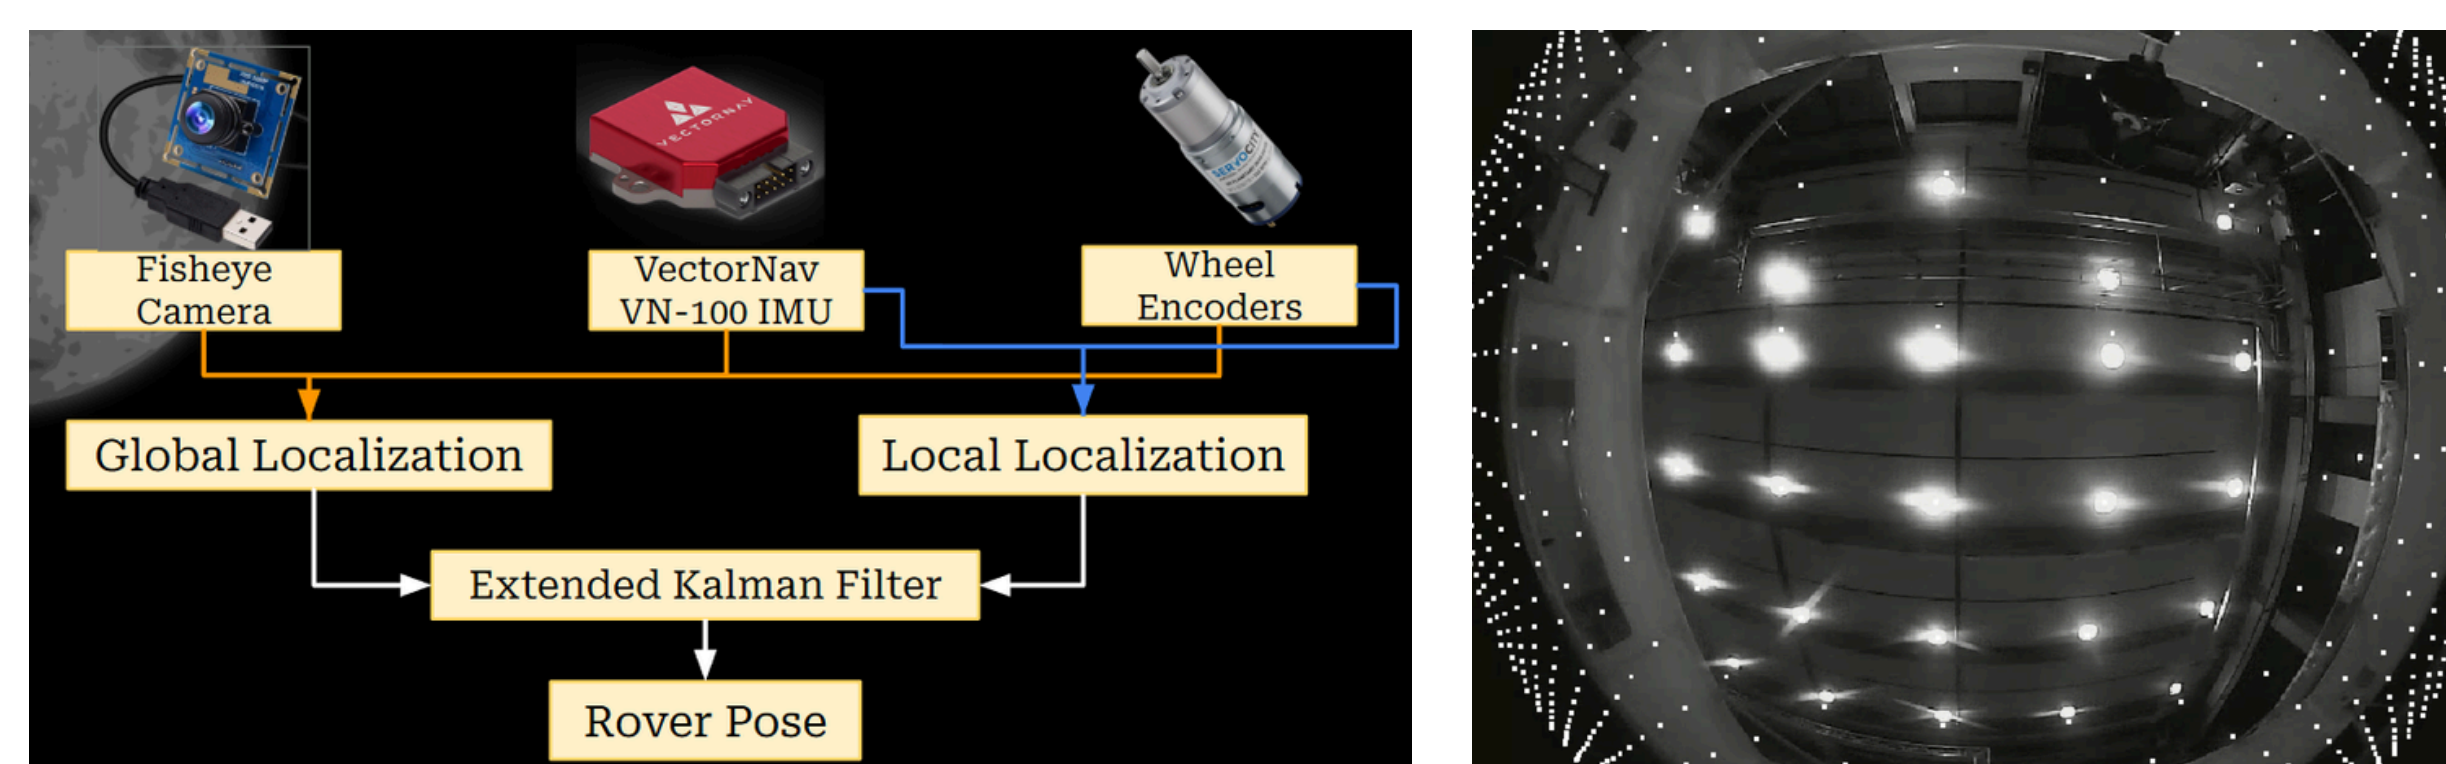
**Crater Detection & Localization:** The rover utilizes RGB-D data streams from a ZED 2i stereo camera to identify terrain features. A YOLOv8-nano model, trained on a custom dataset, performs real-time object detection to locate craters. Upon detection, the system extracts the feature's geometry: the centroid coordinates in the world frame and diameter.

**Adaptive Tool Planning:** The system analyzes the terrain topology before manipulation. A point cloud of the target area is captured and cropped to the specific crater region. The algorithm analyzes this local map to find the maximum elevation peaks. These values are mapped to a calculated Tool Offset, which adjusts the height of the grading blade to match the terrain.

**Post-Grading Validation:** After completing a grading pass, the system autonomously verifies the surface flatness. The raw point cloud undergoes Voxel Downsampling to reduce noise, followed by surface normal estimation. The algorithm calculates individual point slopes and masks out ghost points. Finally, it computes the average slope of the area; if the residual angle is below the target threshold (< 5°), the grading is marked as a success.



### Localization



**Global Localization (Star-Tracking Emulation):** The fisheye functions as a "star tracker," navigating using overhead lights as fixed constellations. The system processes data in five distinct stages:

1. *Get Fisheye Image:* Maps raw pixels to 3D ray vectors using a pre-calibrated lookup table.
2. *Find Bright Pixels:* Isolates overhead lights via thresholding, masking out horizon clutter.
3. *Regress to Sky Grid:* Associates identified features by matching them with the nearest nodes on the known ceiling grid.
4. *Calculate Pixel Offset:* Calculates the residual distance between the observed & expected light positions.
5. *Reproject to 3D:* Minimizes alignment errors to solve for the updated global pose (u, v, θ).

**Local Localization:** For immediate state estimation between visual updates, the system relies on high-speed proprioceptive sensors. **Wheel Encoders** measure distance traveled, while a **VectorNav VN-100 IMU** tracks orientation and rotational velocities. This stream provides smooth, continuous motion data but is naturally susceptible to drift over time.

**Sensor Fusion (Extended Kalman Filter):** The core of the architecture is an Extended Kalman Filter (EKF), which synthesizes these inputs. The EKF uses the local stream to predict the rover's state at high frequency, while utilizing the reprojected 3D data from the vision pipeline to correct that prediction. This ensures the final Rover Pose remains both smooth and globally accurate.

### Navigation & Manipulation

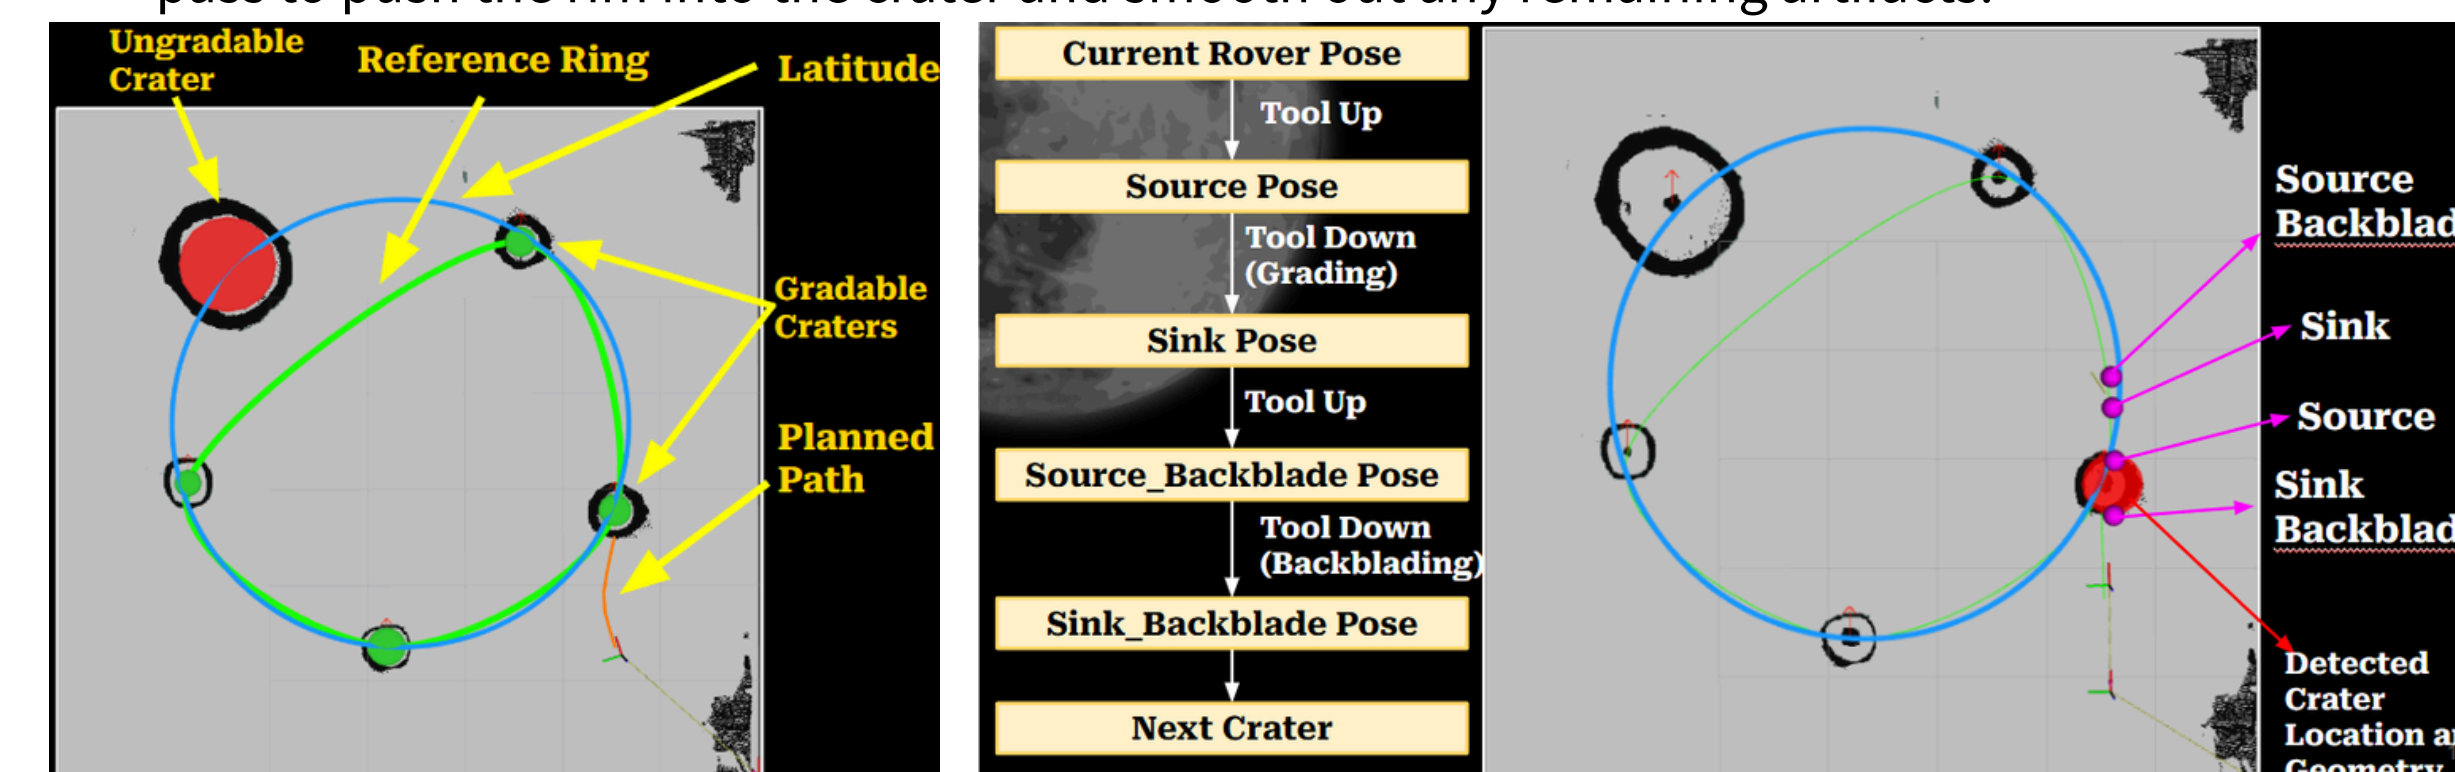
**Global Planning: Lattice A\* with Ackermann Primitives:** The system utilizes a Lattice A\* planner that treats yaw orientation (θ) as a third dimension (x, y, θ). This approach generates smooth, feasible trajectories that are compatible with the rover's Ackermann steering constraints. The global controller is a Pure Pursuit controller, which is also used by the local planner for manipulation.

- *Reference Ring:* The planner generates a bold green spline that acts as a "Reference Ring," derived from a polynomial equation passing through the centroids of all gradable craters.
- *Cost Function:* The criteria enforce Counter-Clockwise (CCW) movement, penalize backward motion during global transit, and minimize cross-track error to "hug" the reference spline.

**Local Planning: Task-Space Maneuvers:** Once the rover approaches a target crater, the system switches to a simple A\* planner for inter-pose navigation. Because the grading vectors are primarily straight lines, this planner explicitly permits bi-directional movement, allowing the rover to execute necessary reverse maneuvers for manipulation.

**Manipulation: Automated Grading Sequence:** The excavation process follows a rigorous state machine designed to move soil and smooth the terrain. As illustrated in the workflow:

1. *Forward Grading:* The rover navigates to the crater, lowers the tool, and drives forward to the crater center, pushing regolith into the crater.
2. *Backblading:* The rover navigates to the opposite end of the crater and performs a backward pass to push the rim into the crater and smooth out any remaining artifacts.



### Results



**Lunar ROADSTER engages full autonomy for end-to-end trail path grooming**

- Grooms craters up to 0.5m in diameter and 0.1m in depth.
- Grooms craters to a mean gradient of 2° (Requirement: < 5°) and validates
- Avoids craters greater than 0.5m in diameter
- Plans a path with a mean deviation of 22.89% from reference latitude (Requirement: < 25%)
- Follows the planned path with a mean deviation of 7.01% (Requirement: < 10%)
- Localizes itself in a GPS-denied environment and communicates to the user
- Has a ground contact pressure of 1.37kPa to prevent sinkage (Requirement: < 1.5kPa)

**Future Work**

- Integration of wheel slip estimation for better localization
- Adding more DOFs to the blade (yaw) for improved manipulation dexterity
- Testing in bigger spaces to identify long-horizon defects
- Automated stuck-state recovery in the Behavior Executive Node to handle failed crater detections without operator intervention