

Received 21 February 2025; revised 19 December 2025; accepted 5 January 2026.

Date of publication 12 January 2026; date of current version 4 February 2026.

This article was recommended by Associate Editor David Wettergreen.

Digital Object Identifier 10.1109/TFR.2026.3652156

# Lunar Rover Cargo Transport: Mission Concept and Field Test

ALEXANDER KRAWCIW<sup>1</sup> (Graduate Student Member, IEEE), NICOLAS OLMEDO<sup>2</sup>,  
FAIZAN REHMATULLAH<sup>2</sup>, MAXIME DESJARDINS-GOULET<sup>3</sup>, PASCAL TOUPIN<sup>3</sup>,  
AND TIMOTHY D. BARFOOT<sup>1</sup> (Fellow, IEEE)

<sup>1</sup>Institute for Aerospace Studies, University of Toronto, Toronto, ON M5S 1A1, Canada<sup>2</sup>MDA Space, Brampton, ON L6Y 6K7, Canada<sup>3</sup>Centre de Technologies Avancées, Sherbrooke, QC J1K 0A5, Canada

CORRESPONDING AUTHOR: ALEXANDER KRAWCIW (alec.krawciw@mail.utoronto.ca)

The work of Alexander Krawciw was supported by the Vanier Canada Graduate Scholarship.

(Field Report)

**ABSTRACT** In future operations on the lunar surface, automated vehicles will be required to transport cargo between known locations. Such vehicles must be able to navigate precisely in safe regions to avoid natural hazards, human-constructed infrastructure, and dangerous dark shadows. Rovers must be able to park their cargo autonomously within a small tolerance to achieve a successful pickup and delivery. In this field test, Lidar Teach and Repeat (LT&R) provides an ideal autonomy solution for transporting cargo in this way. A one-ton path-to-flight rover was driven in a semi-autonomous remote-control mode to create a network of safe paths. Once the route was taught, the rover immediately repeated the entire network of paths autonomously while carrying cargo. The closed-loop performance is accurate enough to align the vehicle with the cargo and pick it up. This field report describes a two-week deployment at the Canadian Space Agency's (CSA) Analog Terrain, culminating in a simulated lunar operation to evaluate the system's capabilities. Successful cargo collection and delivery were demonstrated in harsh environmental conditions.

**INDEX TERMS** Artemis program, lunar operations, rovers.

## I. INTRODUCTION

THE Artemis program is NASA's plan to return astronauts to the Moon [1]. This program has already completed the successful launch of the Orion spacecraft in 2022 [2] and is working toward sending astronauts to the lunar surface by 2027 [3]. Beyond an initial touchdown, there are long-term plans to create a lunar outpost on the moon's surface. Outpost structures will house astronauts and support science on the Moon. A key challenge to supporting such an operation is safely transporting building materials and supplies to assemble and maintain the operation. Uncrewed landers would make deliveries away from the settlement for safety reasons.

This creates a lunar logistics problem: large pieces of cargo need to be reliably transported from the lander sites back to the habitat, and empty cargo must be removed [4]. The long distances involved, currently estimated at 5 km [5], motivate

using an autonomous cargo vehicle, allowing operations to occur without astronauts on the Moon, or freeing them to perform research and science instead of cargo transport.

This article presents a mission concept and associated field test of a lunar logistics mission carrying heavy cargo in a precise manner that enables future operations in the challenging lunar context. This field report takes existing Lidar Teach and Repeat (LT&R) pipelines and evaluates them in a new application domain to evaluate suitability and remaining challenges for future development. This is preparatory Lunar Surface Exploration Initiative work to develop and test technologies and concepts of operation that could enable future Canadian robotic contributions at the lunar surface.

### A. MISSION CONCEPT

To support cargo transport operations on the Moon, we assume that there exists a known set of landing sites,

a known set of cargo delivery locations, and a long distance of relatively unknown terrain between them. Satellite imagery and depth maps can be used in strategic mission planning, but will not contain enough information to allow precise path planning of a rover. In addition, we propose having a dedicated system at the landing site that allows the cargo to be loaded using a crane from a lander onto a transport vehicle. Once the cargo has been loaded, it will be self-sufficient. The rover's requirement is to park precisely relative to a location of interest; the cargo will be able to lift itself off the vehicle to perform the final docking procedures.


These landing sites will be used many times, which means that there is an opportunity to leverage previous driving experiences between the lander and the habitat to speed up autonomous navigation. The teach and repeat framework [6] is ideally suited for this type of mission. The first time the rover navigates a route, it maps the path in a *teach pass*. After just one drive, the rover can autonomously repeat it in either direction along that route. In this concept, the first time the rover travels between the lander, habitat, and any other locations of interest, a network of paths is taught by using short, remotely specified path segments. Another benefit of this exploration approach is the ability to undo if the progress is unsatisfactory. A small section of the recent path can be reversed, and then a new teaching branch can begin in a new direction [7].

In subsequent missions, the network of paths can be reused to repeat from the rover's current location to the location of interest. The critical benefit of a teach-and-repeat approach is that autonomous traversals occur on the same path each time. This means that if the path drives around a boulder or hazard, it will continue to do this every time without having to perform terrain assessment to observe these obstacles again if the environment is static. This reduces computational requirements and allows for more continuous drives between locations of interest, improving energy efficiency and saving time. The only dynamic hazards on the moon will occur in known regions, so operators can move between operational modes depending on the environment.

In this concept, we assume that the cargo is an automated system itself. The cargo connects to the rover to allow the cargo to access rover data and to provide diagnostics back to remote operators through the rover's existing communication channels. Smart cargo provides long-term mission flexibility as new sensors that may be required for docking can be included at the cargo's time of launch rather than upfront with the vehicle.

## B. FIELD TEST

To evaluate this mission concept, the Canadian Space Agency's (CSA) Lunar Exploration Light Rover (LELR) [8] was upgraded with new avionics and outfitted with lidar for situational awareness and compatibility with LT&R [9], [10]. In addition, an actuated cargo was designed to be loaded onto the flat rear bed of the vehicle and carried into position. Fig. 1



**FIGURE 1.** LELR carrying cargo as it drives away from the lander on the right. The cargo transport mission is carried out under simulated lunar south pole lighting conditions, and the ground is illuminated by the vehicle's headlights and the "sun gun." Two Ouster lidars are used for mapping and hazard detection. (Image Credit: CSA.)

shows LELR as it drives away from the lander toward the habitat. The simulated mission executes a small-scale version of the mission concept defined in Section III.

To make the field test as representative of a lunar deployment as possible, the testing was carried out on the Analog Terrain of the CSA [11]. The Analog Terrain is  $120 \times 60$  m and contains sand, large rocks, and larger hill features that challenge rover drive systems and mimic the surface terrains of Mars. The operations were conducted remotely, except for a safety officer who maintained line of sight with the vehicle but was not part of the operational team. The remote operations were bandwidth-limited to 2 Mbps, and a delay of 5 s was inserted in both directions: sending requests to the rover and receiving data back from it. Finally, to recreate the long shadows and hard illumination caused by the sun at the lunar south pole, an 18-kW spotlight, called the sun gun, illuminated the otherwise dark terrain. Some tests were performed under sun-gun illumination and others in complete darkness.

## II. BACKGROUND

Cargo transport will be a paradigm shift for the autonomy of extraterrestrial rovers. Existing rovers such as the Mars Exploration Rovers (Spirit and Opportunity) [12], Curiosity [13], and Perseverance [14] only visit locations once. In missions like these, exploration is the focus. In practice, most missions to date rely on a high level of remote operator involvement for motion planning and hazard avoidance [15]. To explore, rovers need to perceive nearby hazards, plan around them, and estimate their current localized position to achieve strategic mission objectives [16]. In our cargo transport mission concept, exploration capabilities will be required to complete the first drive to new landing sites or habitat delivery locations.

In the static lunar environment, hazard avoidance is performed primarily by assessing the geometric extent of the rover from stereo depth reconstruction [17] and additionally considering the suspension of the vehicle as it moves over 3-D

terrain [18]. Autonomous ground vehicles on Earth often use lidar for terrain assessment off-road [19], [20], but so far no lidar has been sent on an extra-terrestrial rover mission [21].

Strategic planning for exploration is based on the scientific objectives of the mission. Short-range path planning is used to plan short routes that deviate from the strategic plan in the presence of obstacles [22]. Considerations such as slip estimation [23], communication reliability [14], and illumination of solar panels [24] can all be included to find safe and efficient routes. Variations of sampling-based motion planners [25] are often used to select obstacle-free routes, and optimal routes are chosen based on minimizing the energy consumed and the rover-specific considerations above.

Low-drift odometry remains important to cargo vehicle navigation, and the ability to use maps generated in the operational environment opens new avenues for algorithm development and testing. Due to its high reliability, visual stereo odometry is the primary method of relative state estimation used on extraterrestrial rovers [26], [27]. However, if operating in illuminated regions near the lunar south pole, the resulting shadows will cause significant challenges for visual methods because stereo camera mapping quality degrades in shadows [17]. ShadowNav [21] uses a Hapke model of light from an illumination source on the ego-vehicle. This allows the vehicle to localize against known craters in a depth elevation map during lunar night. Fang and ‘Red’ Whittaker [28] use holes in stereo to find shadows and use ray tracing to help determine the relative orientation of the sun to the vehicle, exploiting the presence of shadows as a feature that can be used to improve localization. Lidar odometry methods [29], [30] consistently outperform stereo odometry on standard on-road driving benchmarks [31]. In addition to low drift, lidar methods are invariant to the ambient illuminations and will localize as shadows change. Until a 3-D flight lidar is validated, it is difficult to assess the impact of the solar interference on a lidar at the lunar south pole.

Other field tests have been performed to evaluate cargo transport for a lunar environment. The ATHLETE robot [32] was designed as a rover that could be packaged with the cargo. It features six actuated arms with small wheels at the end of each. Cargo is carried on a central hexagonal platform. On Earth, ATHLETE could carry up to 300 kg of payload [33]. Large cargo cylinders mocking habitat modules were joined into a larger habitat complex in a large-scale field test. The unique actuated arm design enables ATHLETE to move in redundant ways. Under normal operation, the platform rolls on its six wheels and steers holonomically by twisting the arm. However, in case the vehicle gets stuck because of a large rock or soft sand, it can walk by lifting its legs up and down. This minimizes robot mass because smaller wheels and motors can be used. An extension of ATHLETE was proposed to collect and receive payloads. The Tri-ATHLETE concept splits an ATHLETE robot into two half-robots with three arm-legs each. The two vehicles

approach a cargo from either side and latch onto it. The pallet of the cargo forms the central platform and rigidly links the two Tri-ATHLETES into one functioning system. Although the half-rovers are less stable, they can still safely navigate without payloads.

Other, more conventional rover concepts have been proposed for cargo missions. In their full-scale lunar mission analysis, Akin and Bolatto [34] propose using a rover with a U-shaped chassis that can align around standard cargo pallets. In addition to the large transport capabilities of a large rover, they suggest that certain payloads could be attached to a fleet of similar rovers to allow for regolith clearing and site preparation before habitat construction. A fleet of smaller multipurpose rovers would be the only way to achieve the required scale or regolith clearing with planned launch capabilities.

For a cargo mission, the rover will transport loads between known places repeatedly. This motivates the inclusion of maps created on the robot for continued operation in the future. Visual Teach and Repeat (VT&R) uses visual odometry to build up a topometric pose graph that allows relative localization of the robot during autonomous operation. The teach and repeat paradigm is simple: while exploring, capture features from odometry and store them in a local submap. The submaps are joined together topologically with the relative transformations estimated from odometry. These two steps already occur to explore using odometry, so storing the relevant data structure is a lightweight addition to the system. Storing the relative transformations of the robot also encodes the robot’s path. This process of odometry and mapping is the teach phase, and the path the robot drives is called the teach path. Once the robot has driven the teach path, it can immediately drive it again autonomously. Any autonomous traversal is a repeat of the teach path. While repeating, the current sensor input is compared to the map to localize the robot with respect to the teach path. The relative position of the rover to the path forms the error term for a controller that brings the robot back to the teach path. Using only relative information enables robots to repeat the path accurately, even if the odometry used to teach it drifted. In another lunar-simulated field test, VT&R was evaluated on Devon Island, in the high arctic [6]. Over 32 km of driving was completed with an over 99% autonomy rate. A challenge with the stereo camera teach and repeat is the dependence on ambient lighting. Features captured during the day look different from those captured at night under headlight illumination. Extensions of VT&R are robust to illumination changes, either by repeating routes while the illumination changes to capture multiple experiences [35] or by using deep learning to extract feature descriptors that can be matched directly between day and night [36]. Switching to lidar as the primary sensor [10] prevents issues related to ambient illumination. This work is an application of the lidar odometry and localization developed by Burnett et al. [10] and the model-predictive control developed by Qiao et al. [9].

The Lunar Light Exploration Rover was initially designed to support uncrewed exploration or manual piloting from astronauts onboard [37]. Its significant payload capacity allowed flexibility in adding science instruments or human pilots. VT&R [6] was demonstrated as an autonomy solution that allowed the rover to use stereo imagery to follow a mapped route. Using a combination of the visual odometry from VT&R with local and global planning, LELR completed a large-scale test campaign in 2015 [38].


For the Artemis mission, a large volume of lunar cargo will be transported from landers to the Artemis base. NASA estimates a cargo demand of up to 10 000 kg per year [5]. Currently, planned cargo includes habitation modules, communication systems, and the logistics infrastructure. In addition to one-time items required to establish operations, resupply missions bring food, water, and spare parts. All of this cargo will require different surface mobility solutions to move goods to their appropriate final destination. High cargo volumes motivate the use of autonomous transport vehicles on the lunar surface.

### III. MISSION PLAN

The field test is designed to evaluate the viability of the proposed mission concept and expose new technical challenges for further development. At a high level, there are three stages of the tests: a semi-autonomous stage where a teach is performed, an autonomous stage where the mission is completed in reverse, and a forward autonomous stage that repeats all of the mission motions completed in the teach. In a lunar deployment, the manual control stage would only be performed once to complete the long traverse between the habitat and lander sites. This could be any mixture of teleoperation, astronaut control, or a slower autonomy stack. Autonomous repeats can occur any number of times after the first teach. Even if the landers change positions, a repeat can be used to bring a rover near the landing site, and then another operating mode can perform the fine alignment with the lander.

Fig. 2 shows the high-level route that was selected from satellite and elevation images. The apex is a location close to the habitat where the rover will switch direction to align. The lander and habitat require the rover to align within a prescribed tolerance to ensure successful cargo operations. The semi-autonomous mission stage has seven steps as follows.

- 1) Drive from the start to the lander semi-autonomously and align with the lander.
- 2) Transfer cargo from the lander to the vehicle using a crane.
- 3) Drive forward from the lander to the apex semi-autonomously.
- 4) Drive backward from the apex to the habitat, aligning the cargo.
- 5) Unlatch and raise the cargo from the vehicle bed.
- 6) Drive the vehicle forward so that it is clear of the cargo.
- 7) Lower the cargo so that it is completely aligned.



**FIGURE 2. Key locations for the simulated lunar mission. The lander and habitat require precise relative alignment. The other locations are approximate and will be selected based on an operator's assessment of navigability. (Image Credit: CSA.)**

At this point, the rover has demonstrated a successful delivery. The second mission stage essentially resets the environment by performing the entire mission in reverse. The most precise motion is the second operation: reversing the vehicle back under the raised cargo. To latch successfully, the rover requires positioning of  $\pm 7.5$  cm longitudinally and  $\pm 10$  cm laterally. In addition, the cargo legs have  $\pm 15$  cm of clearance from the centerline, and there is a 20-cm space between the vehicle's cab and the cargo. The mission reset stage has seven opposite steps as follows.

- 1) Raise the cargo so that the rover can park underneath.
- 2) Reverse the vehicle under the cargo.
- 3) Lower the cargo onto the vehicle bed and latch it.
- 4) Drive forward to the apex autonomously using a repeat.
- 5) Drive in reverse from the apex to the lander using a repeat.
- 6) Unlatch the cargo and move it from the vehicle to the lander using the crane.
- 7) Drive in reverse from the lander to the start.

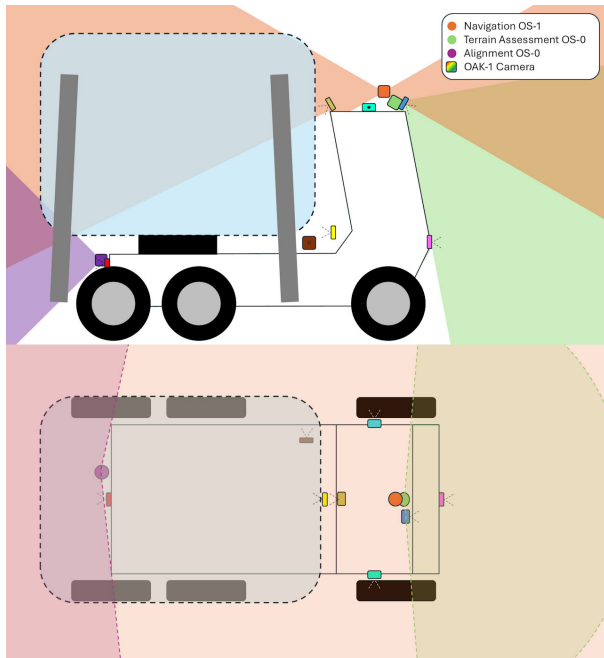
After resetting the environment, the first mission stage is executed again with autonomous driving for all motions. The lander and habitat positions are recorded precisely by the teach path, so operators simply instruct the vehicle to repeat them. The autonomous mission has four steps as follows.

- 1) Drive forward to the lander autonomously using a repeat.
- 2) Load the cargo using the crane.
- 3) Drive forward to the apex and then reverse align to the habitat autonomously using repeat.
- 4) Deploy the cargo.

Autonomous missions should occur more quickly than semi-autonomous missions. To compare the two missions, the completion time, path consistency, and cargo handling were assessed.

### IV. HARDWARE

LELR is a six-wheel drive vehicle with two independently steered wheels controlled by electric actuators and four fixed




**FIGURE 3.** Side and top view showing the locations of the three lidar sensors and eight cameras. The lidar fields of view are denoted by the swept area filled with each sensor's color. Note that the cargo (*light blue*) reduces the navigation lidar's (*orange*) field of view as the vehicle carries it. The direction of each camera is marked with a dashed triangle.

wheels. The steering can replicate an Ackermann geometry virtually. The rover possesses the capability for zero-radius turning through skid-steering, achieved by independently modulating the rotational velocities of the left and right wheel sets via dedicated electric wheel motors. This allows for differential wheel speeds, including counter-rotation, enabling highly maneuverable, on-the-spot pivoting. However, under typical operational conditions, the rover employs a steering mechanism analogous to that of a conventional automobile, prioritizing controlled, Ackermann-like steering dynamics to optimize stability, minimize tire wear, and reduce energy consumption on standard terrains. The motion of the rover is controlled by a low-level dynamics controller that converts body velocity commands from high-level control loops into the torques for the six wheels and angles for the steering actuators [39].

The rubber tires have a diameter of 0.763 m, allowing obstacles up to 30 cm high to be cleared. LELR is powered by a set of lithium-ion batteries with a total capacity of 12 kWh at 96 V. The vehicle has an unloaded mass of 800 kg and can carry up to 300 kg at a maximum velocity of 13 km/h. The front cab was added to improve the viewpoint of the primary sensors, position lights, and emulate the radiators and solar panel arrays required for energy transfer on the Moon.

Remote communications used a 9-dBi antenna at 2.4 GHz. The link has an allowable range of approximately 300 m from the base station. The ground control station antenna is placed sufficiently far away from the GCS building to reduce interference.



**FIGURE 4.** Composite view of all eight cameras mounted to LELR. The border colors correspond to the camera's color in Fig. 3.

## A. SENSORS

LELR has three lidars, eight cameras, an inertial navigation system (INS), and position feedback on the cargo latches. The primary navigation sensor is an Ouster OS-1 128 lidar [40] mounted parallel to the body and centered on top of the cab. This lidar has the largest field of view, which is beneficial while mapping. The rear third of the field of view is obstructed when the cargo is loaded, so the view of the environment changes depending on the cargo configuration. An Ouster OS-0 128 is tilted toward the ground and mounted on the top of the cab. This lidar provides a view right in front of the vehicle, filling in the near-range blind spot of the OS-1 and providing terrain information for navigability assessment. The third lidar is an Ouster OS-0 128 mounted below the cargo on the rear of the vehicle. The rear view is required to operate LELR semi-autonomously while driving in reverse. Self-occlusions obstruct the ground near the vehicle even when there is no cargo. This lidar captures the habitat precisely as the vehicle reverses to align the cargo.

The eight cameras improve the situational awareness of remote operators but are not currently integrated into the autonomy stack. The front and rear deck cameras show the terrain along the path of the rover. Four cameras are mounted on the cab roof and tilt down to show the footprint of the vehicle. Two additional cameras are mounted on the cargo deck for operational support. One camera views the cargo alignment guide to monitor cargo loading. The other looks sideways at the lander as the rover drives alongside it. Fig. 4

shows a view from all eight cameras after the cargo has been loaded onto the vehicle.

A Novatel Span G370 INS is mounted to the vehicle's roof and provides ground truth for evaluating path tracking. GPS postprocessing was used to improve the precision of the INS pose estimates with Novatel's Inertial Explorer software.

Remote cargo handling requires sensor feedback to ensure that each step is completed safely and correctly. The actuators that clamp the cargo onto the vehicle's bed have position and current feedback. Position feedback is used when releasing the cargo to ensure that the hooks are clear of the latching rod. Current feedback is used when latching the cargo as a proxy for force feedback. The cargo is securely clamped when the actuator current exceeds a calibrated threshold.

While the cargo is lifting, an IMU inside it provides feedback on roll and pitch. This information is used to adjust the power for each leg and keep the cargo safe and level for the final docking.

After the rover has driven clear of the cargo, the legs retract to achieve the target alignment height. A downward waterproof ultrasonic distance sensor measures the height of the cargo off the ground.

## B. CARGO

The cargo module measures 1.6 m in diameter and 2.2 m in length. It is constructed from expanded polystyrene, which optimizes both structural integrity and weight efficiency. The interior of the cargo is hollow, providing accommodation for various electrical subsystems. The door frame facilitating access to these components was 3-D printed to ensure precision and structural adequacy in a single, seamless piece. The complete cargo assembly has a total mass of 182 kg, balancing the requirements for structural resilience, weight constraints, and operational functionality.


### 1) TELESCOPIC LEG MECHANISM

The cargo is supported by four independently actuated telescopic legs, each powered by dedicated electric actuators to allow precise deployment and retraction. A simplified plastic roller system facilitates smooth telescopic motion, with tolerances and openings designed to prevent small debris from obstructing movement. Each leg allows 1.52 m of travel, allowing it to retract within the rover's footprint.

When deploying the cargo, the legs begin fully retracted. The legs extend to elevate the cargo up to 0.5 m above the deck, allowing it to support itself independently of the rover. With the cargo securely self-supported, the rover disengages and departs, leaving the cargo in its designated location. After the vehicle has departed, the legs retract for final alignment with the habitat.

### 2) ALIGNMENT AND LOAD DISTRIBUTION SYSTEMS

The alignment guides feature 45° tapered surfaces, allowing for tolerance of misalignments up to  $\pm 150$  mm between the rover deck and the cargo. The contact surfaces of these guides are covered with 3.175-mm ultrahigh-molecular-



**FIGURE 5. Three-frame sequence of a latch release. When the latch is fully engaged, as in the first pane, the cargo is secure, and it is safe to drive the rover. When the latch is clear, loading and unloading operations can occur. (Image Credit: CSA.)**


weight polyethylene sheets, which allow the cargo to slide into the correct position to be latched. The cargo latches are shown in Fig. 5.

The legs pivot with respect to the cargo body as they retract or extend. The leg end caps are designed to inhibit outward slipping while allowing controlled inward motion, improving the system's stability.

## V. AUTONOMY STACK

The backbone of the autonomy used for this cargo mission is LT&R [10]. LT&R has been shown to be effective for long-term navigation in harsh and unstructured environments such as forests, grassy plains, and deep snow [9], [41]. The method follows the same framework as VT&R [6], but lidar offers a few advantages, especially for lunar deployments. Lidar sensors are invariant to ambient lighting. This means that even with the low elevation of the sun at the Moon's south pole, the sensor would not be blinded. Long shadows that change as the Moon orbits make long-term feature matching more difficult for stereo camera localization [42].

The teach phase uses lidar odometry to estimate the relative transformations between nearby submaps along the route. The Ouster OS-1 lidar is used for odometry and mapping



**FIGURE 6.** Lidar map of the CSA's Analog Terrain created by accumulating the lidar scans from LELR as it completed the demonstration in simulated lunar south pole lighting conditions. Points are colored by their elevation relative to the start position. The rover's estimated path is shown in black.

because it has the largest field of view around the vehicle. Our lidar odometry uses a robust variant of the iterative closest point (ICP) algorithm [10]. Point-to-point associations form the loss function, but the cost weighting is based on the estimated normal of the map, providing many of the optimization benefits of a point-to-plane approach. Adding robust cost functions rejects outlier points that do not match. This improves localization performance with scans that include the cargo matching to maps that do not contain it. The odometry drift was evaluated as 0.54% on the Boreas dataset [10] using the Kitti metric, which is lower than visual odometry [43]. The odometry drift is small enough to accumulate useful maps while operating. Fig. 6 is an accumulation of one scan every 3 m along the route, transformed using the odometry roll-out to that location. Qualitatively, the edge of the habitat (top middle) and other distinct features like craters are well defined.

While teaching new routes, a pose is stored every 30 cm or 10 degrees of motion. Submaps are stored less frequently. The submaps are accumulations of recent scans, cropped to 40 m around the rover, and voxel downsampled. Accumulating scans means that blind spots can be filled, and map points are distributed more evenly. Uniform map density improves localization if a different lidar configuration is used. The operator requests branch and merge points to create an arbitrary network of paths. During a repeat, the vehicle may take any route through the branches regardless of the order in which they were taught.

The repeat phase allows the robot to autonomously navigate the route after it has been driven. Lidar odometry is computed first, and the compounded transformation of odometry and the prior localization estimate is used as a prior in the localization optimization. When driving autonomously,

the submap nearest to the robot is selected, and ICP is used to estimate the robot's position within that submap. As the robot moves along the route, new submaps for localization are selected. The robust cost function in ICP improves localization: dynamic obstacles in either the map or live scan are rejected as outliers.

The lidar odometry is accurate enough to use dead reckoning to close the control loop for distances a few meters long without an appreciable impact on the path tracking. While the full odometry and map localization pipeline usually runs in real-time (i.e., greater than 10 Hz), map localization may take longer depending on the vehicle's motion and the map's quality. Instead of reducing localization quality to maintain speed, the exponential moving average of the runtime is used to decide on the fly whether to skip map localization. This ensures that the controller will receive state updates at 10 Hz and allows ICP map localization to take the appropriate number of iterations to converge.

Once localization is complete, the relative cross-tracking error between the vehicle and the taught path forms the input to a model-predictive-control (MPC) problem. The teach path defines the target robot poses over a finite horizon. The MPC aims to follow the target path at an appropriate speed and uses a unicycle model centered between the four rear wheels to determine an optimal control command [9]. A speed scheduler adjusts the target forward velocity based on the curvature of the path. The command is a planar twist: forward speed and angular velocity. The MPC solves for commands at 15 time steps spaced 0.25 s apart. The first command from the solution is sent to the vehicle. With the 0.1-s time for each lidar scan, an updated MPC problem will be solved before the prediction window is complete. A low-level controller in the vehicle converts these twist commands

into motor torques for all six wheels, and a steering angle for the front two [39]. To reduce the MPC's modeling error, a first-order lag is added to the angular velocity command.

## VI. SEMI-AUTONOMOUS DRIVING


Direct teleoperation of LELR was deemed infeasible due to the limits placed on sensor bandwidth and the 10-s round-trip delay between the ground control station and the vehicle. A semi-autonomous driving paradigm was implemented to provide a mode of exploration based on human terrain assessment. An operation begins with the approximate global location of the vehicle, based on lidar odometry estimates and other large features visible in satellite data. Sensor data is triggered by the driver and sent back to the ground control station. Depending on the direction of motion, the appropriate cameras and lidar are queried. After receiving the scans, the rover driver creates a short path that is safe for the rover to traverse. To define the path, the driver uses a joystick to move a 3-D model of the vehicle through the 3-D point cloud map accumulated by the lidar. The model contains the same kinematic constraints as the real vehicle, ensuring that the selected path is valid and smooth. If the operator is unsatisfied with the path rollout, they can remove part of it and try again or restart completely. Some parts of the mission are best defined using delta pose commands rather than virtual driving. For example, when driving out from underneath the raised cargo, the rover must drive straight to prevent a collision with the cargo legs. In this case, a change in the longitudinal position of the rover can be specified using a delta position.

The paths always begin from the rover's current estimate of its location, which prevents aggressive path recovery behaviors at the beginning of motions. Once the driver is satisfied with the visualization of the future robot motion, shown in Fig. 7, the rover tracks the path based on the rollout of lidar odometry measurements. The low drift rate of lidar odometry leads to accurate motions that reflect the desired paths.

When the rover approaches a location where precise alignment is required, for example, the habitat, relative measurements between the rover and the target are used to guide the next motions. At the lander, the vehicle must align parallel to the cargo for a successful loading operation. The side deck camera has an alignment guide overlay to determine when the vehicle is in position. At the habitat, the rear lidar determines the distance between the vehicle and the docking port. The rear camera also has alignment overlays on the legs of the habitat. Alignment guides can be created based on known camera parameters and 3-D models of the structures of interest. The final manual alignment of the vehicle is shown at both the lander and habitat in Fig. 8(a) and (b), respectively.

## VII. EXPERIMENTS


The field test evaluation was performed after sunset from 5:45 to 10:30 P.M. The temperature started at  $-3\text{ }^{\circ}\text{C}$  at 6 P.M. and fell to  $-6\text{ }^{\circ}\text{C}$  at 11 P.M. [44]. More significantly, the wind speed was 46 km/h at 6 P.M. and gradually dropped to 24 km/h at 11 P.M. [44]. Due to strong winds, it was unsafe to lift the




**FIGURE 7.** Final target rover pose (translucent gray) at the end of the semi-autonomous path segment (red). The red line denotes the target path remaining, and the green line tracks the actual motion of the real rover (black). The point cloud map is an accumulation of scans stitched together using lidar odometry. The real rover will stop moving at the pose of the target rover. The blue lines show the current arc of the front and rear wheels based on the current joystick steering angle.

cargo with the crane three times as originally proposed in the mission plan. The cargo lift was performed only once during the semi-autonomous mission stage.

Fig. 9(a)–(h) shows the semi-autonomous mission stage as a sequence of frames. Initially, the rover is empty at the start [Fig. 9(a)] and drives until it is aligned with the lander [Fig. 9(b)]. The cargo was loaded using the telehandler and supported with four human-guided tow lines. In Fig. 10, the crane operation is performed. The rover can track the operation using the  $360^{\circ}$  view from the OS-1 lidar. The lidar and rear deck camera allowed operators to measure precisely that the cargo was seated correctly and ready to be latched. The high winds necessitated four tow lines to prevent the cargo from swinging. Even in earlier testing with lower wind speeds, the cargo was difficult to carry with the telehandler. Once the cargo is loaded onto the vehicle bed, it is secured with a pair of latches. The hook closes into the cargo and pulls down until it is tight. In Fig. 5, the unlatching sequence is shown. With the cargo loaded, the rover drives the bulk of the route [Fig. 9(c)] and switches directions at the apex [Fig. 9(d)]. The alignment into the habitat occurs with LELR driving in reverse. Fig. 9(e) shows the cargo during the final approach. Once the cargo is aligned, it can be released. First,



(a)



(b)

**FIGURE 8.** Alignment guides overlaid with the static structures on the lander (a) and habitat (b). When the green lines cover the legs of the structures, operators know that the vehicle is aligned.

the cargo is unlatched, as shown in Fig. 5. Next, the legs extend [Fig. 9(f)] until all four are touching the ground. As they extend, the cargo rises above the cargo deck until the state of Fig. 9(g). Once clear, LELR drives forward to leave the cargo behind [Fig. 9(h)]. The final height adjustment of the cargo occurs, and the legs retract again to set the cargo to the desired final position.

Three types of performance were evaluated: alignment accuracy with the lander and habitat, path-tracking error, and mission time.

The rover's alignment quality is presented in Table 1. The distance measurements were taken with a tape measure by the safety officer once the flight team declared that an operation was complete. When aligned with the lander, the rover should be parallel and between 1.5 and 2.0 m laterally away. Because the telehandler moves the cargo laterally from the lander to the rover, there is a wider margin for error. The heading should be within  $\pm 5^\circ$  of parallel. The front of LELR was 1.8 m, and the rear was 1.7 m, both within the distance tolerance. The heading misalignment was  $2.7^\circ$ , which is sufficiently parallel for the alignment guides to work properly.

The alignment with the habitat targeted less than 0.15 m between the cargo and the habitat without impact. Laterally, the requirement is  $\pm 0.05$  m for the centreline of the cargo with the habitat. The final position of the cargo was 0.17 m away from the habitat and 0.01 m laterally offset. The longitudinal gap was slightly too long, but the cargo was laterally aligned well.

GPS was used to determine the ground truth of the rover's position. Fig. 11 shows the route taken over the terrain. The control point of the path-tracking controller is the center of the rear four wheels. The path-tracking errors are summarized in Table 2. The LT&R estimate at the rear is the controller error. This RMSE is 0.10 m. Applying the known extrinsic calibration between the rear axle and the GPS generates the LT&R estimates at the GPS. The rms path-tracking error at the GPS is 0.142 m, and the maximum is 0.472 m. This error

**TABLE 1.** Alignment results for the rover during the semi-autonomous teach mission phase.


Alignment Target	Value	Result
Distance from Lander	1.5 m - 2.0 m	1.75 m
Parallel Angle to Lander	$0 \pm 5^\circ$	$2.7^\circ$
Longitudinal Gap from Habitat	(0, 0.15] m	0.17 m
Lateral Misalignment from Habitat	$0 \pm 0.05$ m	0.01 m

**TABLE 2.** Path-tracking performance while performing the cargo mission. All values in meters.

Description	Mean $\pm \sigma$	RMSE	Max Error
LT&R estimate at rear	$0.024 \pm 0.112$	0.101	0.401
LT&R estimate at GPS	$0.015 \pm 0.167$	0.142	0.472
GPS measurement	$-0.041 \pm 0.248$	0.252	0.573

is visualized in Fig. 12. The Ackerman nature of this vehicle means that heading errors at the rear axle lead to larger lateral errors at the front of the vehicle. The GPS ground-truth path-tracking rms error at the front of the rover is 0.252 m. The maximum error is 0.573 m. This is larger than the LT&R estimate in the same location, suggesting that errors are a mixture of controller errors and localization errors. Comparing the relative positions of the teach and repeat data from the GPS and the relative transform from ICP provides an RMSE of 0.208 m in translation and 0.036 rad. This localization error is higher than previous works [10] but is small enough for the rover to track the path accurately.

The semi-autonomous teach was built from 18 path segments that were followed by the rover. In open stretches, the longest path was 15.92 m long. In the alignment steps, the shortest path was 0.33 m. The total path length is 164.2 m. These segments are visualized with different colors in Fig. 13. There were no e-stop interventions activated. The operators' terrain assessment and path execution kept the rover away from hazards. Near the habitat, the extra motion occurs due



(a)



(b)



(c)



(d)



(e)



(f)



(g)



(h)

**FIGURE 9.** Composite panel showing the snapshots of the rover as it completes the semi-autonomous mission plan. (a) Rover at the start. (b) Rover aligned with the lander. (c) Rover rounding the south curve. (d) Rover at the apex. (e) Cargo aligned with the habitat. (f) Cargo legs extended. (g) Cargo lifted off the rover. (h) Rover drives away from the cargo. Image Credit: CSA.



**FIGURE 10.** Cargo is lifted using the telehandler and loaded onto LELR. The top row is a static camera, and the lower row is the synchronized view from the rover's OS-1 lidar. The high winds required the use of extra personnel with tow lines.




**FIGURE 11.** Ground-truth route of the rover superimposed on the satellite view of the Analog Terrain.

to the vehicle not tracking a maximum steering angle turn precisely. This led the operators to direct the rover to move out and align again.

The timing breakdown for the mission is compared between semi-autonomous control and autonomous control in Table 3. In total, driving semi-autonomously took 2:02:14. The same distance of autonomous driving took 10:32, only 8.6% of the time. The cargo lift time could be sped up with a fully automatic procedure. Driving out from under the cargo was always performed autonomously. This short repeat took 55 s.

### VIII. DISCUSSION

Many observations from our field test require further investigation or consideration in future designs. The quality of the lidar localization suggests that most of the path-tracking error is a result of the path-tracking controller. The most significant model mismatch occurs from the MPC assuming that the rear of the robot moves as a unicycle. While this assumption is fairly accurate for shallow curves, the response



**FIGURE 12.** Absolute path-tracking error of LT&R evaluated at the position of the GPS, near the front of the vehicle. The vehicle steers back and forth across the taught path. The largest errors occur near the sharpest turns.

of the Ackermann vehicle was difficult to model accurately. In addition, the tuning of the controller occurred with different inertial parameters. The cargo was not yet ready for use, so 180 kg of cylindrical weights were added to the rear deck. The cargo has a much larger inertia because of the larger volume. While the paths are not the same, LT&R path-tracking error



**FIGURE 13.** Eighteen path segments were executed during the semi-autonomous driving portion of the mission. These segments were recorded to create the smooth teaching path that was repeated afterward. Note that segments are short during precise alignments.

**TABLE 3.** Duration of mission steps for both forward missions. The speed-up of continuous motion is clear in the precise alignment with the habitat.

Mission Step	Semi-Autonomous Duration (mm:ss)	Autonomous Duration (mm:ss)
Drive Start to Lander	25:20	2:42
Load Cargo	16:54	N/A
Drive Lander to Apex	41:55	5:44
Align from Apex to Habitat	56:59	2:06
Deliver Cargo	18:00	N/A

estimates from the earlier tests were 0.06-m RMSE, which is significantly smaller than the 0.10 m observed during this demo. Future controllers should consider whether or not the cargo is loaded and adapt accordingly. Wang et al. [45] employ adaptive MPC control for an autonomous forklift that has dynamics that change with the type of cargo moved. The low-level torque controller only considered payloads weighing up to 80 kg [39], which will compound the problem.

Another issue with the unicycle assumption is the implicit symmetry of forward and reverse motion. Ackermann vehicles in reverse are nonminimum phase dynamic and much

more difficult to control. Using this vehicle meant that the most precise part of the mission must occur with LELR driving in a mode that is more difficult to control. Operationally, we mitigated this by providing a long straight route to the habitat, which gave the controller more time to reduce the lateral error. This extra direction switch occurs between the apex and Habitat in Fig. 11. A symmetrical vehicle would be easier to control, but considering a bicycle model that accounts for reverse Ackermann control issues should mitigate the issue in future deployments.

Transferring the cargo from the lander to the vehicle using a crane was challenging. The load was supported by one strap, which left it free to swing and twist with respect to the crane. Strong winds during the final demo exacerbated the cargo motion, but even on earlier test days that were relatively calm, the telehandler alone could prevent the cargo from swinging out of position and colliding with the vehicle’s cab. The pendulum motion of hanging cargo will be significantly more pronounced on the lunar surface with gravity 1/6th of the Earth’s. Cargo transfer was not the focus of this field test, but lunar solutions that involve cranes will require significant support to constrain the motion of large, heavy items. Gill et al. [46] investigate a number of cargo offloading concepts. They note that an important consideration for cargo offloading is whether landers would integrate a payload-specific offloading solution or a general-purpose solution, like the mobile crane used here, would be more cost-effective. There is a tradeoff between heavier cargo that is self-contained and lighter cargo that requires extra infrastructure. Observations of this field deployment suggest that alternatives to a mobile crane should be considered.

At both locations where precise alignment was required, further automation would improve the accuracy and decrease the time. Automatic alignment would allow more flexibility for landers in various positions. The time it took to make small, manually computed adjustments to the rover’s position near the lander and habitat was significantly longer than the average rate of driving. This occurred for a few reasons: first, due to the nonholonomic constraints of the vehicle, correcting lateral errors required backing out a significant distance and trying again. Second, sensing blind spots made it more difficult to ascertain exactly where the parts at risk of a collision were located, and third, the error margin of path completion was not tunable. While aligning with the habitat, the drive team decided that the lateral error was greater than 10 cm and out of spec on the first approach. To correct this, a new alignment approach would be performed. Because every motion was being taught, the drive team requested that the robot repeat back along the path 5 m from the habitat. Once the robot had completed the repeat, a new teach branch was created with the second successful alignment.

Part of the challenge with alignment was the occlusions of useful sensors. At the lander, the rear lander leg was blocked in the OS-1 lidar scan by the roof lights on the cab, which made it more difficult to assess whether or not the vehicle was parallel. Once the cargo was loaded, the lander

was completely blocked from lidar scanning. At the habitat, the rear camera and lidar gave a clear view of the reverse approach. However, the cargo should be instrumented with a distance sensor at the docking location. It was unreliable to assess the longitudinal position of the cargo on LELR once it was loaded. This created uncertainty of  $\pm 7.5$  cm on the location of the cargo. While reversing, a conservative envelope had to be used even though the cargo was ultimately loaded almost at the center. This resulted in being 17 cm away from the habitat rather than less than 15 cm as was the target. A direct measurement of the alignment parameters would improve teleoperations and allow for a fully automatic approach to be developed.

The last capability to improve alignment is a tunable speed profile for the rover. The current path-tracker has the same target and limit speeds independent of the path length. For short motions less than 30 cm, the behavior needs to change. First, the rover should move more slowly to prevent damage in case the alignment is suboptimal. Second, the tolerance to decide whether the end of the path has been reached should decrease. For long motions, it was satisfactory to use  $\pm 5$  cm as a complete path threshold, but for a motion that might only be 10 cm, this did not work. The drive team added offsets to the desired motion to account for the lower level behavior, which should be avoided.

The wireless communication between the rover and the ground station was mostly reliable. However, there were occasional signal dropouts. The omnidirectional antenna used at the ground station has two stronger lobes. Aligning one stronger lobe with the most common location of the rover improved the stability. Longer autonomous repeats were not affected directly by communication dropouts because no operator interaction was required. One future extension of using LT&R while exploring would be a communication recovery behavior where the rover could backtrack along the path to a location with a reliable communication link.

## IX. LESSONS LEARNED

Performing a two-week deployment provided valuable insights into the mechanical system design, software user interfaces, and operational procedures.

Performing field testing in the Canadian winter meant that snow was anticipated, but the weather variability was underappreciated in project planning. Throughout the deployment, the weather proceeded from dry and cold, to 10 cm of snowfall, to strong winds, to 24 h of rain, which melted the snow, and to scattered ice patches when the temperature dropped below freezing again. The external sensors were IP-rated and could withstand snowfall, but the weather pushed the system beyond some of its original design specifications. Traction became a variable that was unaccounted for. When the sand was dry, the rubber tires left clear tread marks and displaced the sand while driving. When the snow fell, the traction was reduced, but with packed snow, the rover still drove well. But when the wet ground froze over, the rover had trouble navigating as reliably as before. The ground was frozen, and

the tires left almost no marks in the sand as the rover moved. In addition, localized wheel slip was visible when watching the rover. The rover never lost traction overall, but the low-level controller was less able to produce the desired body motions.

Another challenge was the site safety conditions for safety operators walking around the test site. Snow clearing and salting efforts were taken to prepare the site, but the amount of human labor required was not anticipated. The rover safety officer handling the emergency stop ensures that the rover and test site are not damaged if something goes wrong. When the temperature is below freezing and windy, proper warming gear is critical to allow testing to proceed. Large-button e-stops, remote controls, and walkie-talkies had to be implemented so that operators would not have to remove their gloves to activate the robot. While it might seem convenient to schedule field deployments for summer, it is useful to challenge the system's robustness even if these weather variations are not directly applicable to the Moon. In addition, the earlier sunset during the winter provides more testing time under lunar illumination.

One interesting result was the robustness of the lidar localization to snow accumulation. Teach passes were completed to build maps before the snow fell. LELR could easily localize its pose and then follow the original path even though the snow obscured many small geometric features, such as rocks. With the amount of appearance change, a human operator could not have followed the original path.

Using teach and repeat as an operator motivates several useful back-end changes to better support cargo missions. Previous deployments of teach and repeat do not contain tight alignment stages in the middle of a route. Even when driving with the joystick, an alignment attempt would be made while teaching, but the final position or path was deemed unsatisfactory. If a teach is continued while backing up and driving forward to align again, all subsequent repeats would include the erroneous alignment and direction switch. In this field trial, we did a short repeat back to an upstream location along the path, branched onto a new teach, and tried the alignment again. This achieved the desired effect as future repeats all ignore the original stub branch. However, this created an operational problem as the two routes would be so close to each other that the visualization would often overlap. Without a layer selection, it was difficult to ensure that a target waypoint was on the correct branch. This motivates a new feature to mark branches as stale so they do not appear in the user interface as locations for waypoint selection.

Another useful feature during alignment would be a "bookmark position" button that forces a pose-graph vertex to be exactly in the current rover location. To reduce storage requirements and smooth high-frequency oscillations caused by odometry, LT&R only stores a pose every 30 cm or  $10^\circ$ . The ideal aligned location is unlikely to sit exactly 30 cm away, especially in a place such as the lander, which is not at the end of the path. A workaround was developed by stopping and restarting the teach in the target location. This

is undesirable because the two teach segments are no longer locally consistent through an odometry rollout and lead to tiny direction switches or jumps in the target path.

The final operator feature to add is an offset from the waypoint for the end of paths. When LELR aligned autonomously to the habitat, the teach took the cargo until it was less than 5 cm from the contact interface. While LT&R could achieve this precision, out of caution, there was interest in making the initial alignment have a gap and then proceeding. The smallest gap that could be requested was 30 cm, which did not bring the rover acceptably close. Adding the functionality to go to the end minus 10 cm would have been ideal.

## X. CONCLUSION

Through an extensive two-week deployment of the LELR, a successful cargo transport concept was demonstrated and evaluated. The mission concept considered the case of moving cargo between the same lander and habitat repeatedly. LT&R was demonstrated as an effective framework for precise navigation in a lunar analog environment. Cargo alignment was within 20 cm of contact, path tracking errors were about 25-cm RMSE, and repeating took only 8.6% of the time of semi-autonomous motion. Future extensions of the concept include automatic alignment of the rover to the habitat and lander. Additional monitoring of the route may be required to ensure that the taught path remains stable. Change detection techniques could be used to monitor the effects of the terrain, looking for new obstacles or growing trenches from the rover's repeated motion. These improvements would improve the resilience of the system to changes in the lunar environment caused by the Artemis program. This capability would unlock functionality that supports cargo landers that do not land in the same spot every time.

## ACKNOWLEDGMENT

The authors thank the team at the Canadian Space Agency (CSA), David Gingras, Fiona Kirby, Cody Barth, Martin Picard, and Tongxi Wu, for their support with the Analog Terrain and operations. In addition, they thank Igor Coelho Correia for taking many of the detailed pictures used in this article. They also thank Hexagon for the use of Inertial Explorer.

This work was made possible through the CSA Lunar Surface Exploration Initiative and an NSERC Alliance Grant.

## REFERENCES

- [1] NASA. (2024). *Artemis*. [Online]. Available: <https://www.nasa.gov/humans-in-space/artemis/>
- [2] NASA. (2022). *Artemis*. [Online]. Available: <https://www.nasa.gov/mission/artemis-ii/>
- [3] NASA. (2025). *Artemis*. [Online]. Available: <https://www.nasa.gov/mission/artemis-iii/>
- [4] *Lunar Mobility Drivers and Needs*, National Aeronautics and Space Administration (NASA), Washington, DC, USA, 2024.
- [5] *Lunar Surface Cargo*, NASA, Washington, DC, USA, 2024.
- [6] P. Furgale and T. D. Barfoot, "Visual teach and repeat for long-range rover autonomy," *J. Field Robot.*, vol. 27, no. 5, pp. 534–560, Sep. 2010.
- [7] B. E. Stenning, C. McManus, and T. D. Barfoot, "Planning using a network of reusable paths: A physical embodiment of a rapidly exploring random tree," *J. Field Robot.*, vol. 30, no. 6, pp. 916–950, Nov. 2013.
- [8] Canadian Space Agency. (2015). *Lunar Light Exploration Rover*. [Online]. Available: <https://www.asc-csa.gc.ca/eng/rovers/rover-fleet/lelr.asp>
- [9] X. Qiao, A. Krawciw, S. Lilge, and T. D. Barfoot, "Radar teach and repeat: Architecture and initial field testing," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2025, pp. 13021–13027.
- [10] K. Burnett, Y. Wu, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Are we ready for radar to replace LiDAR in all-weather mapping and localization?" *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 10328–10335, Oct. 2022.
- [11] Canadian Space Agency. (2021). *Analogue Terrain*. [Online]. Available: <https://www.asc-csa.gc.ca/eng/laboratories-and-warehouse/analogue-terrain.asp>
- [12] M. Maimone, A. Johnson, Y. Cheng, R. G. Willson, and L. Matthies, "Autonomous navigation results from the Mars exploration rover (MER) mission," in *Proc. 9th Int. Symp. Experim. Robot.*, 2006, pp. 3–13.
- [13] A. Rankin, M. Maimone, J. Biesiadecki, N. Patel, D. Levine, and O. Toupet, "Mars curiosity rover mobility trends during the first 7 years," *J. Field Robot.*, vol. 38, no. 5, pp. 759–800, Aug. 2021.
- [14] V. Verma et al., "Autonomous robotics is driving perseverance rover's progress on Mars," *Sci. Robot.*, vol. 8, no. 80, p. 3099, 2023.
- [15] A. Connell and M. Hurst, "Ground software to support autonomous onboard scheduling for Mars perseverance rover," in *Proc. IEEE Aerosp. Conf.*, Mar. 2023, pp. 1–10.
- [16] C. Wong, E. Yang, X.-T. Yan, and D. Gu, "Adaptive and intelligent navigation of autonomous planetary rovers—A survey," in *Proc. NASA/ESA Conf. Adapt. Hardw. Syst. (AHS)*, Jul. 2017, pp. 237–244.
- [17] L. Gerdes, M. Azkarate, J. R. Sánchez-Ibáñez, L. Joudrier, and C. J. Perez-del-Pulgar, "Efficient autonomous navigation for planetary rovers with limited resources," *J. Field Robot.*, vol. 37, no. 7, pp. 1153–1170, Oct. 2020.
- [18] K. Otsu, G. Matheron, S. Ghosh, O. Toupet, and M. Ono, "Fast approximate clearance evaluation for rovers with articulated suspension systems," *J. Field Robot.*, vol. 37, no. 5, pp. 768–785, Aug. 2020.
- [19] P. Krüsi, P. Furgale, M. Bosse, and R. Siegwart, "Driving on point clouds: Motion planning, trajectory optimization, and terrain assessment in generic nonplanar environments," *J. Field Robot.*, vol. 34, no. 5, pp. 940–984, Aug. 2017.
- [20] Y. Shan et al., "LiDAR based traversable regions identification method for off-road UGV driving," *IEEE Trans. Intell. Vehicles*, vol. 9, no. 2, pp. 3544–3557, Feb. 2024.
- [21] A. Cauligi et al., "ShadowNav: Crater-based localization for nighttime and permanently shadowed region lunar navigation," in *Proc. IEEE Aerosp. Conf.*, Mar. 2023, pp. 1–12.
- [22] M. Sutoh, M. Otsuki, S. Wakabayashi, T. Hoshino, and T. Hashimoto, "The right path: Comprehensive path planning for lunar exploration rovers," *IEEE Robot. Autom. Mag.*, vol. 22, no. 1, pp. 22–33, Mar. 2015.
- [23] H. Inotsume, T. Kubota, and D. Wettergreen, "Robust path planning for slope traversing under uncertainty in slip prediction," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3390–3397, Apr. 2020.
- [24] N. Otten, D. Wettergreen, and W. Whittaker, "Strategic autonomy for reducing risk of sun-synchronous lunar polar exploration," *Field Service Robot.*, vol. 5, pp. 465–479, Sep. 2017.
- [25] A. Orthey, C. Chamzas, and L. E. Kavrakı, "Sampling-based motion planning: A comparative review," *Annu. Rev. Control, Robot., Auto. Syst.*, vol. 7, no. 1, pp. 285–310, Jul. 2024.
- [26] M. Maimone, Y. Cheng, and L. Matthies, "Two years of visual odometry on the Mars exploration rovers," *J. Field Robot.*, vol. 24, no. 3, pp. 169–186, Mar. 2007.
- [27] V. Verma et al., "Enabling long & precise drives for the perseverance Mars rover via onboard global localization," in *Proc. IEEE Aerosp. Conf.*, Mar. 2024, pp. 1–18.
- [28] E. Fang and W. 'Red' Whittaker, "Ray tracing and use of shadows as features for determining location in lunar polar terrain," in *Proc. Field Service Robot.*, 2021, pp. 307–320.
- [29] P. Dellenbach, J. Deschaut, B. Jacquet, and F. Goulette, "CT-ICP: Real-time elastic LiDAR odometry with loop closure," in *Proc. Int. Conf. Robot. Autom.*, 2022, pp. 5580–5586.

- [30] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "FAST-LIO2: Fast direct LiDAR-inertial odometry," *IEEE Trans. Robot.*, vol. 38, no. 4, pp. 2053–2073, Aug. 2022.
- [31] L. R. Agostinho, N. M. Ricardo, M. I. Pereira, A. Hiolle, and A. M. Pinto, "A practical survey on visual odometry for autonomous driving in challenging scenarios and conditions," *IEEE Access*, vol. 10, pp. 72182–72205, 2022.
- [32] B. H. Wilcox, "ATHLETE: A cargo-handling vehicle for solar system exploration," in *Proc. Aerosp. Conf.*, Mar. 2011, pp. 1–8.
- [33] B. H. Wilcox, "ATHLETE: A cargo and habitat transporter for the moon," in *Proc. IEEE Aerosp. Conf.*, Mar. 2009, pp. 1–7.
- [34] D. L. Akin and N. Bolatto, "Lunar base construction, operations, logistics, and maintenance: Lessons from microgravity and academia," in *Proc. AIAA SCITECH Forum*, Jan. 2025. [Online]. Available: <https://arc.aiaa.org/doi/abs/10.2514/6.2025-1862>
- [35] M. Paton, K. MacTavish, M. Warren, and T. D. Barfoot, "Bridging the appearance gap: Multi-experience localization for long-term visual teach and repeat," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2016, pp. 1918–1925.
- [36] Y. Chen and T. D. Barfoot, "Self-supervised feature learning for long-term metric visual localization," *IEEE Robot. Autom. Lett.*, vol. 8, no. 2, pp. 472–479, Feb. 2023.
- [37] R. McCoubrey et al., "A Canadian lunar exploration light rover prototype," in *Proc. International Symposium on Artificial Intelligence, Robotics, Automation in Space*, 2012, pp. 1–5.
- [38] J. N. Bakambu, C. Langley, K. C. Leineweber, R. Mukherji, and T. D. Barfoot, "Maturing Canadian autonomous guidance, navigation, and control of planetary rovers," in *Proc. Int. Symp. Artif. Intell., Robot. Autom. Space, i-SAIRAS*, Beijing, China, 2016, pp. 19–22.
- [39] R. Chhabra, "Dynamical reduction and output-tracking control of the lunar exploration light rover (LELR)," in *Proc. IEEE Aerosp. Conf.*, Mar. 2016, pp. 1–8.
- [40] Ouster OS1 LiDAR. Accessed: Feb. 1, 2025. [Online]. Available: <https://ouster.com/products/scanning-lidar/os1-sensor>
- [41] J. Schn, T. D. Barfoot, and J. Collier, "Off the beaten track: Laterally weighted motion planning for local obstacle avoidance," *IEEE Trans. Field Robot.*, vol. 1, pp. 249–275, 2024.
- [42] M. Allan et al., "Planetary rover simulation for lunar exploration missions," in *Proc. IEEE Aerosp. Conf.*, Mar. 2019, pp. 1–19.
- [43] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 3354–3361.
- [44] Weather Underground. (2024). *Saint-Hubert Airport, Quebec, Canada Weather History*. [Online]. Available: <https://www.wunderground.com/history/daily/ca/saint-hubert/CYHU/date/2024-12-12>
- [45] Y. Wang, K. Sun, W. Zhang, and X. Jin, "A velocity-adaptive MPC-based path tracking method for heavy-duty forklift AGVs," *Machines*, vol. 12, no. 8, p. 558, Aug. 2024.
- [46] T. R. Gill et al. (2022). *Lunar Surface Cargo Offloading Concepts*. [Online]. Available: <https://ntrs.nasa.gov/citations/20220005086>

...