

Madrid, Spain

May 5th-7th

2026

uc3m

Universidad
Carlos III
de Madrid

AIAA

Enhanced Sensor Fusion for Autonomous Navigation in GPS-Denied Unknown Environments of Space Robots

Niccolò Goretti

MSc student in Space Engineering, Department of Aerospace Science and Technology (DAER), Politecnico di Milano , Milano, Italy.

niccolo.goretti@mail.polimi.it

Annachiara Ippolito

PhD student in Aerospace Engineering, Department of Aerospace Science and Technology (DAER), Politecnico di Milano , Milano, Italy.

annachiara.ippolito@polimi.it

Mauro Massari

Associate Professor of Space Systems, Department of Aerospace Science and Technology (DAER), Politecnico di Milano , Milano, Italy.

mauro.massari@polimi.it

ABSTRACT

Exploring unknown, GPS-denied environments poses critical challenges, especially in space missions where autonomous navigation is essential. Simultaneous Localization and Mapping (SLAM) enables Unmanned Ground Vehicles (UGVs) to navigate autonomously in such conditions, but reliable localization and mapping must overcome sensor noise, environmental uncertainty, and drift. This work presents a GPS-free, multi-sensor SLAM system for space UGVs that integrates data from a Light Detection and Ranging (LiDAR) sensor, a Monocular Camera (MC), and an Inertial Measurement Unit (IMU) through enhanced sensor fusion. A modified Graph-Based (GB) method fuses MC and IMU localization estimates to reduce drift, while an Extended Kalman Filter (EKF) processes LiDAR-based localization in a second fusion stage, refining the trajectory. Simulations conducted in a virtual environment validate the proposed method under realistic conditions. The results show significant reductions in positional drift, improved localization robustness, and superior performance compared to a standard odometry-based EKF. By combining complementary sensors in a modular framework, the approach improves SLAM for autonomous navigation in unstructured terrains and demonstrates strong potential for planetary exploration and other GPS-denied missions.

Keywords: Autonomous Navigation, GPS-denied Environments, Space Rovers, Unmanned Ground Vehicles, Sensor Fusion, Simultaneous Localization and Mapping

Nomenclature

δ_{rot1}	=	Initial rotation command (odometry model)
δ_{trans}	=	Translation command (odometry model)
δ_{rot2}	=	Final rotation command (odometry model)
δ_{rot1_n}	=	Noisy initial rotation command
δ_{trans_n}	=	Noisy translation command



$\delta_{\text{rot}2_n}$	=	Noisy final rotation command
$\epsilon_{\text{rot}1}, \epsilon_{\text{rot}2}, \epsilon_{\text{trans}}$	=	Gaussian noise terms for control commands
x, y	=	Position in Absolute Reference Frame
θ	=	Heading angle (yaw about vertical axis) in the planar formulation
s	=	Along-track displacement in Relative Reference Frame
n	=	Cross-track displacement in Relative Reference Frame
h	=	Vertical axis in Relative Reference Frame
ψ, θ, ϕ	=	Euler angles in the 3D formulation (yaw, pitch, roll)
p	=	State / pose vector (x, y, θ)
\hat{p}	=	Optimized pose (Graph-Based fusion)
p_{IMU}	=	Pose estimate from Inertial Measurement Unit
p_{MC}	=	Pose estimate from Monocular Camera
W_{IMU}	=	Confidence weight matrix for Inertial Measurement Unit
W_{MC}	=	Confidence weight matrix for Monocular Camera
$J(p)$	=	Local weighted least-squares cost function used in GB fusion
P	=	State covariance matrix
Q	=	Process noise covariance matrix
R	=	Measurement noise covariance matrix
p_{pred}	=	Predicted state (Extended Kalman Filter)
P_{pred}	=	Predicted covariance (Extended Kalman Filter)
z	=	LiDAR measurement vector
v	=	Measurement noise vector in the EKF observation model
H	=	Observation model matrix
I	=	Identity matrix
K	=	Kalman gain
p_{est}	=	Estimated state after Extended Kalman Filter update
σ	=	Standard deviation
$\overline{\text{err}}_s$	=	Mean along-track error
$\overline{\text{err}}_n$	=	Mean cross-track error
$\overline{\text{err}}_\theta$	=	Mean heading error

1 Introduction

Space exploration continues to drive scientific and technological advances, with the Moon serving as a key target and testbed for deep-space missions. Lunar subsurface structures such as lava tubes and impact crater caves offer natural shelters from radiation, micrometeoroids, and thermal extremes, while preserving valuable geological records [1–5]. However, autonomous navigation in these environments faces significant challenges due to the lack of global references (GPS, magnetic fields, lighting), as well as unknown and unmapped terrain, irregular features, dust, temperature gradients, communication delays, and limited computational resources [6–8]. These factors underscore the need for efficient real-time localization and mapping solutions.

To address these challenges, this work proposes a multi-sensor Simultaneous Localization and Mapping (SLAM) framework for unstructured GNSS-denied settings, commonly referred to as GPS-denied. While motivated by lunar cave exploration, it is designed to remain applicable across a broad range of scenarios. The approach fuses data from an Inertial Measurement Unit (IMU), a Light Detection and Ranging (LiDAR), and a Monocular Camera (MC) to exploit their complementary strengths. A hybrid wheeled-legged Unmanned Ground Vehicle (UGV) is chosen to enhance terrain adaptability, with a two-layer sensor fusion architecture combining a modified graph-based (GB) method and an Extended Kalman Filter (EKF) to improve stability and accuracy. Validation is performed in a high-fidelity

simulation pipeline using Blender™, Unreal Engine™, MATLAB™, and Simulink™. The study focuses on planar motion as a first-step simplification, aimed at isolating the core localization and sensor-fusion issues while maintaining a scenario representative of constrained ground navigation in lunar subsurface contexts. This assumption is consistent with the flat simulation setup considered in this work and provides a controlled baseline for future extensions toward full 6-DoF motion on uneven terrain.

SLAM, a core tool for GPS-denied navigation, has evolved from probabilistic filtering approaches to graph-based and hybrid formulations [9–11]. Early methods based on EKF provide a compact and recursive framework for joint state estimation, and have been widely adopted in robotics applications [10]. GB approaches, on the other hand, improve global consistency by formulating SLAM as an optimization problem over a set of poses and constraints [11]. These methods typically rely on loop closures or feature-based associations to correct accumulated drift, achieving high precision in structured environments. Hybrid solutions combining filtering and optimization techniques have been proposed to balance accuracy and computational efficiency [12], while the need for consistency in pose and landmark estimation has been emphasized since early work [13]. Despite these advances, many SLAM implementations still rely predominantly on a single sensing modality or on tightly coupled assumptions between sensing and motion models. This can limit system resilience in unstructured and feature-scarce settings, such as lunar subsurface scenarios, where odometry may be degraded by terrain interaction and feature extraction may be affected by poor illumination and repetitive geometries [12, 14]. These considerations motivate the adoption of multi-sensor fusion strategies, where complementary sensing modalities are combined to improve robustness and reduce reliance on any single source of information in demanding operational contexts.

Hardware and sensor choices critically impact navigation performance. Common alternatives, such as dead reckoning, beacon-based, time-of-flight (ToF), landmark-based, and vision-based methods, each have drawbacks in subsurface scenarios [15]. Multi-sensor fusion combining visual, inertial, and ToF-based sensing has shown superior results [16, 17]. Accordingly, this work adopts an IMU, LiDAR, and MC with artificial lighting to support stable perception. The chosen sensor suite reflects sensing configurations increasingly explored in planetary robotics, where inertial sensing, vision-based perception, and active range measurements are combined to improve robustness in challenging environments [8, 18]. IMUs provide high-rate motion estimation but are affected by drift, monocular cameras enable feature-based localization when sufficient visual information is available, and LiDAR offers accurate range measurements that are independent of ambient illumination. This complementarity makes the selected combination particularly suitable for subsurface and low-visibility scenarios, such as lunar caves, where passive sensing alone may be insufficient [19]. Furthermore, the considered sensors are compatible with typical space robotics constraints, as ongoing developments are progressively improving their size, weight, and power characteristics for robotic exploration applications [8, 18].

Mobility is addressed by selecting a hybrid wheeled-legged UGV, balancing terrain adaptability and efficiency [20]. While UAV-UGV cooperative systems are effective in terrestrial GPS-denied exploration, lunar aerial mobility demands alternative propulsion such as cold gas thrusters due to the lack of atmosphere, which shapes the proposed multi-platform architecture.

To better position the method with respect to conventional SLAM approaches, it is worth highlighting that the framework departs from standard control-driven EKF formulations by introducing a fusion-driven prediction strategy, as well as from landmark-based methods by relying on direct sensor pose estimates. Within this context, the main contributions of this work can be summarized as follows:

- A two-layer loosely coupled SLAM architecture in which a graph-based fusion stage is used to generate an optimized trajectory that replaces traditional control-based prediction in the EKF, reducing the impact of odometry noise and drift.

- A landmark-free graph-based fusion strategy that combines IMU and MC pose estimates at each time step, specifically designed for feature-scarce and GPS-denied environments such as lunar subsurface scenarios.
- A fusion framework that operates independently of explicit control inputs, enabling robust trajectory estimation even in the presence of actuator uncertainty and partial sensor unavailability.
- A quantitative assessment of the proposed approach against a standard odometry-driven EKF, demonstrating improved accuracy in along-track, cross-track, and heading estimation.
- A modular simulation pipeline integrating Unreal Engine, MATLAB, and Simulink to support realistic validation of multi-sensor SLAM strategies in controlled settings.

This paper is structured to progressively build upon the presented approach, guiding the reader through the methodology, experimental results, and concluding discussions. Section 2 details the proposed methodology, including the system modular configuration and the two-layer sensor fusion framework. Section 3 outlines the simulation workflow and presents the corresponding results after each processing stage. It then compares the method against a standard EKF based on odometry and concludes with a sensitivity analysis on the EKF noise parameters. Section 4 concludes the paper with a summary of findings and future research directions.

2 Methods

The chosen methodology addresses the key challenges of autonomous navigation in GPS-denied environments through a modular and scalable multi-sensor SLAM framework. Designed to ensure real-time operation, robustness, and adaptability, it integrates data from an IMU, a LiDAR, and a MC within a loosely coupled architecture. Validation is performed through a high-fidelity simulation pipeline that combines photorealistic rendering, physics-based sensor emulation, and algorithmic processing. A concise summary of the hardware configuration and the modular structure developed for data integration and processing is presented here. With this foundation established, the focus of this section shifts to a detailed description of the navigation system design and operational flow.

2.1 Reference Scenario Definition

The reference scenario builds upon the key findings discussed earlier in the introduction, serving as a foundation for introducing the navigation system underlying the considered SLAM framework.

2.1.1 Hardware Setup

The proposed system is composed of a deployer, a UAV, and a UGV, each contributing to localization, communication, exploration, and environment sensing. The deployer functions as a reference point, thus as a landmark, and as a communication hub, equipped with dual radio transceivers that act as beacons and enable coordination across platforms. It also supports logistical functions such as sample storage, power supply, and environmental monitoring. The UAV enhances mapping through high-resolution imaging and atmospheric sensing and may also assist localization by deploying beacons. The UGV, equipped with a hybrid wheeled-legged configuration, autonomously navigates using SLAM, performs terrain sensing, and collects soil and rock samples. Each component is outfitted with a specific sensor suite: the deployer includes dual radio transceivers, the UAV carries a Stereo Camera (SC), an IMU, and a transceiver, and the UGV is equipped with MC, IMU, LiDAR, and a radio transceiver, in line with sensing combinations currently adopted or actively investigated for robotic planetary navigation and cave exploration [19, 21, 22]. This sensor distribution supports robust, cooperative aerial-ground exploration in GPS-denied environments.

2.1.2 Navigation System Architecture

The architecture allowing real-time localization, mapping, and navigation is illustrated in Fig. 1. It integrates data from all platforms to ensure continuity and adaptability and is composed of six modules:

- **Data Interface Module (DIM):** Handles the acquisition and preprocessing of raw sensor data, ensuring storage and compatibility.
- **Sensor Fusion Module (SFM):** Integrates multiple sensor inputs onboard the UGV to estimate its local and global pose while generating a preliminary environmental map.
- **Localization Module (LM):** Uses the UGV's local pose from SFM and refines it through triangulation with radio beacons, improving accuracy.
- **Mapping Module (MM):** Enhances the initial map coming from the SFM by incorporating aerial data from the UAV's SC.
- **Path Planning Module (PPM):** Determines optimal navigation routes based on localization and mapping data, dynamically adjusting paths in response to environmental changes.
- **Control Module (CM):** Converts planned trajectories from PPM into precise actuator commands, guiding the UGV along its path while compensating for control noise.

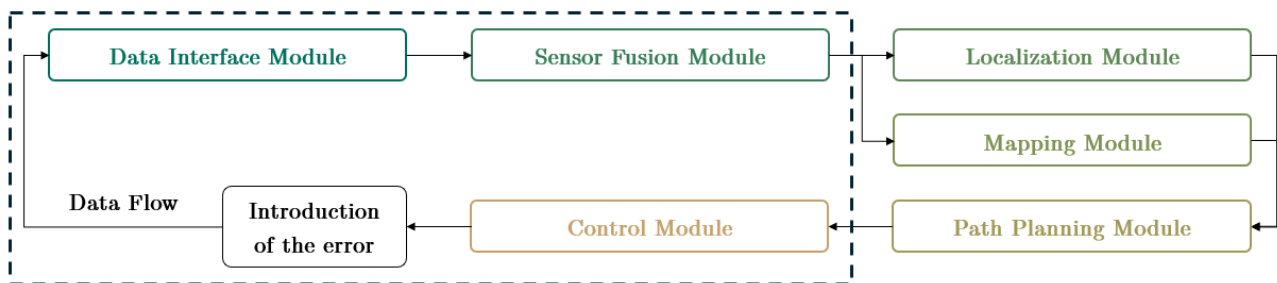


Fig. 1 Modular architecture designed for the navigation system: data flow and module interactions (developed modules in the dashed box)

A key design concern is control noise in CM, due to actuator imprecision. To reduce reliance on noisy odometry, the system emphasizes advanced fusion techniques using environmental data. This structure ensures modularity, fault-tolerance, and resilience in GPS-denied environments.

As anticipated, the core focus of this paper is the development of an effective SLAM method tailored for a UGV. With reference to Fig. 1, the implementation concentrates on the modules supporting the early stages of navigation, specifically the DIM, the SFM, and the CM, including the modeling of control error. Therefore, the following paragraphs focus on the detailed design of these key modules.

Concerning the PPM, which dynamically determines the most efficient navigation routes, it is used in this work solely to provide control inputs as constraints for the SLAM process and is not developed in detail. Similarly, the LM, which integrates the UGV's local pose from the SFM with triangulation data from radio beacons to enhance localization precision, and the MM, which refines the preliminary environmental map from the SFM with aerial data from the UAV's SC, are treated as downstream components. While integral to the system, the detailed implementation of LM and MM is beyond the scope of this paper.

2.2 Design Breakdown

With the navigation architecture and sensor integration strategy well-defined, the focus now shifts to the structured workflow adopted for implementing the navigation algorithm. This section begins by describing the key modules designed for this work, starting from the control commands that are provided to the UGV. These commands govern its movement, allowing the initiation of sensor data acquisition. Subsequently, the focus shifts to the fusion of sensor data with the aim of achieving optimal localization

accuracy. Finally, a brief explanation of the fundamental reference frames used in the simulation is introduced, serving as preparation for the detailed setup description and ensuring clarity across the various stages of the workflow.

2.2.1 CM - Control Commands and Error

The CM serves as the actuator interface of the system, issuing movement commands to the UGV based on navigation inputs. In this study, no external inputs are provided to the CM. Instead, both the trajectory and the associated control commands are derived from a predefined path simulating the expected output of the PPM. These signals are manually defined to reflect the nominal system behavior and are expressed as odometry-based motion instructions, including initial rotation (δ_{rot1}), translation (δ_{trans}), and final rotation (δ_{rot2}), computed from the desired trajectory using a standard odometry motion model [9].

To emulate realistic execution of motion commands, actuation uncertainties are introduced directly at the control input level. These uncertainties are not intended to model sensor noise, which is treated separately in the perception modules, but rather to capture the typical discrepancies between commanded and executed motion observed in real robotic systems. In practice, mobile robots rely on onboard odometric feedback (e.g., wheel encoders) to estimate their motion. Instead of explicitly modelling such sensors and their error sources, this work considers a simplified approach in which their net effect is represented as perturbations on the control commands. This allows reproducing realistic deviations from the nominal trajectory without introducing an additional layer of sensor modelling.

The adopted noise model follows the standard probabilistic odometry formulation described in [9], where control inputs are affected by zero-mean Gaussian perturbations applied independently to translational and rotational components. Accordingly, the translation and rotational components (δ_{trans} , δ_{rot1} , δ_{rot2}) are perturbed to simulate execution errors arising from actuation imperfections such as wheel slip, terrain interaction variability, and control inaccuracies. Formally, the noisy control commands are generated as:

$$\begin{aligned} \delta_{\text{trans}_n} &= \delta_{\text{trans}} + \epsilon_{\text{trans}}, & \delta_{\text{rot1}_n} &= \delta_{\text{rot1}} + \epsilon_{\text{rot1}}, \\ \delta_{\text{rot2}_n} &= \delta_{\text{rot2}} + \epsilon_{\text{rot2}} \end{aligned} \quad (1)$$

where each ϵ term is drawn from a zero-mean Gaussian distribution. The standard deviation is set to 0.01 m for translational motion and 0.01 rad for rotational components, corresponding to approximately 5-10% of typical command magnitudes along the considered trajectories. This choice reflects a realistic level of actuation uncertainty commonly assumed in mobile robotics applications, particularly in unstructured environments such as planetary surfaces, while remaining independent from system-specific calibration.

As a result, the CM outputs two sets of control commands: an ideal set, representing perfect execution of the planned trajectory, and a perturbed set, representing a more realistic scenario affected by actuation errors. These control inputs are then used to drive the UGV in the simulation, enabling the comparison between nominal and realistic motion conditions for subsequent processing by the navigation filter.

2.2.2 DIM - Raw Data Collection

The DIM acts as the data acquisition gateway of the system, collecting raw measurements from the UGV's onboard sensors. It operates in parallel with the CM, receiving both ideal and perturbed control inputs to initiate simulation. The ideal control inputs serve to initialize the UGV's movement, allowing the sensors to start capturing data along the predefined trajectory. In contrast, the noisy controls are used as a baseline for traditional EKF processing, providing a reference against which the performance of the proposed SLAM approach can later be assessed.

During the simulation, the DIM captures and organizes measurements from the sensor suite, including LiDAR point clouds, MC images, and IMU data. These raw datasets act as inputs for subsequent

localization and mapping. Specifically, the LiDAR provides dense 3D spatial information about the environment, the MC delivers high-resolution visual frames, and the IMU supplies measurements of linear acceleration and angular velocity that reflect the rover’s dynamics. The data acquisition process implemented in the DIM ensures that each sensor’s data stream is properly captured and stored, ready for use in the SFM.

2.2.3 SFM - Navigation Filter Outline

The navigation filter is implemented in the SFM using a multi-layer loose coupling approach, illustrated in Fig. 2. This design enables each sensor to operate individually while producing separate pose estimates that are selectively fused. Compared to tightly coupled methods, this structure simplifies integration, troubleshooting, and modular adaptation, supporting the replacement or modification of sensor modules without extensive system adjustments [23].

In the first filtering layer, pose estimates are generated from the three primary sensors. At this stage, each one operates autonomously, ensuring modularity while establishing a solid foundation for data fusion.

The second filtering layer refines sensor data through a two-loop fusion process. In the first loop, the IMU and MC pose estimates are fused using a modified GB method that employs a weighted approach to optimize their respective contributions. In the second loop, an EKF refines the unified estimate obtained from the first loop by integrating the LiDAR pose estimate, producing the final fused trajectory with improved accuracy.

Overall, the SFM receives as input the raw data streams collected by the DIM, comprising LiDAR point clouds, MC images, and IMU measurements. These data enable the computation of independent pose estimates for each sensor, which are subsequently fused through the described process. The module outputs the fused pose estimate and a preliminary map of the environment generated from LiDAR data. These outputs provide essential inputs for the downstream LM and MM, which use them as a baseline.

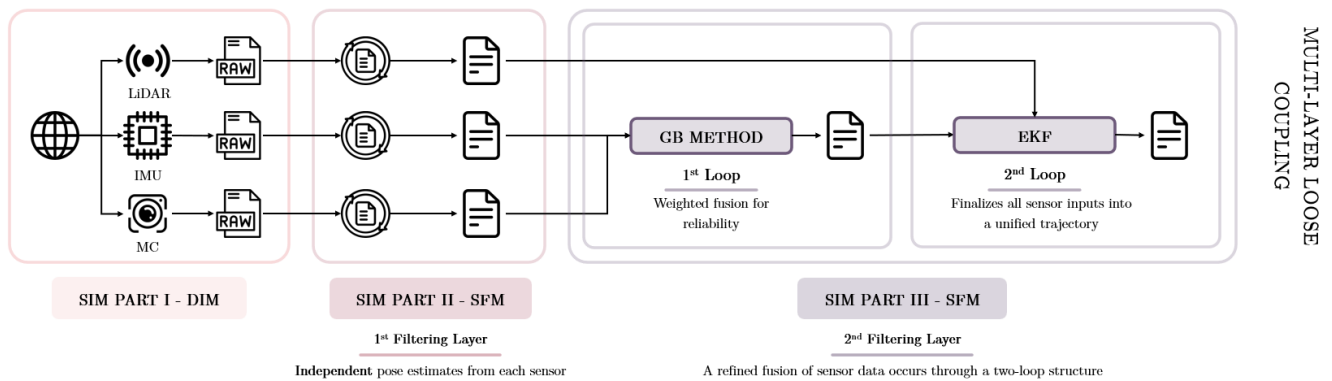


Fig. 2 Navigation filter outline: multi-layer loose coupling

2.2.4 Reference Frames

The simulation framework uses four reference frames, each with a specific role:

- **Absolute Reference Frame (ARF):** This frame is defined in the context of the Automated Driving Toolbox [24]. Its conventions follow the official documentation, which describes the Earth-fixed reference frame (x, y, z) as inertial and adhering to the right-hand rule. The axes are defined such that x points in the vehicle’s initial forward direction, parallel to the ground plane; y is orthogonal to x , also parallel to the ground plane; and z points upward, perpendicular to the ground plane.
- **World Coordinate System (WCS):** Unreal Engine™ uses a world coordinate system (X, Y, Z) that is also inertial but follows the left-hand rule. Its origin is located at the center of the

simulation environment. Although its axes are oriented similarly to the ARF, the difference in handedness necessitates a transformation. To ensure consistency when integrating data from MATLAB™/Simulink™ into Unreal Engine™, it is necessary to invert the y -axis and all rotations around it.

- **Relative Reference Frame (RRF):** The vehicle-fixed coordinate system, denoted as (s, n, h) , is non-inertial and moves with the rover. The axes are defined as s (Along-Track), representing the forward direction; n (Cross-Track), the leftward direction, perpendicular to s ; and h (Height), the upward direction. The associated rotational degrees of freedom are ψ (Yaw), which is the rotation about the h -axis; θ (Pitch), the rotation about the n -axis; and ϕ (Roll), the rotation about the s -axis.
- **Pattern Reference Frame (PRF):** A camera-relative reference frame (X_P, Y_P, Z_P) is used for MC data. The nomenclature used mirrors the fact that, to estimate the parameters of a MC, calibration is typically performed using multiple images of a calibration pattern, such as a checkerboard. In this frame, X_P points to the right in the image plane; Y_P points downward in the image plane; and Z_P points outward from the image plane, meaning towards the camera. Since the MC is onboard the rover, its reference frame is inherently linked to the vehicle's motion.

As anticipated in the introduction, the present study focuses on planar motion in order to reduce problem dimensionality and isolate the main localization and fusion mechanisms. This choice is consistent with the adopted simulation environment, which is intentionally flat and does not include significant terrain elevation changes or attitude variations. Accordingly, rover motion is described only in terms of in-plane translations and heading variation, while vertical displacement and out-of-plane rotations (roll and pitch) are neglected. The vertical axis is therefore retained solely as the reference for heading definition. This assumption aligns with the concept of a 2D pose, a widely used representation in robotics and SLAM to describe a vehicle's state in terms of position and orientation on a plane. The resulting 2D pose is expressed in two different reference frames:

- **Absolute Pose:** Defined in the ARF as (x, y, θ) , where x, y represent the vehicle's absolute position, and θ represents the heading angle, defined as the yaw rotation about the h -axis of the RRF.
- **Body-Centric Pose:** Defined in the RRF as (s, n, θ) , where s is the forward displacement (Along-Track), n the lateral displacement (Cross-Track), and θ the yaw angle about the vertical axis.

The absolute pose (x, y, θ) is primarily used for trajectory estimation and visualization, assuming an initial condition of $(x, y) = (0, 0)$ and $\theta = 0$. This establishes a global reference for position and heading as the rover navigates through the environment. However, to assess deviations from the intended trajectory, a transformation to the body-centric pose (s, n, θ) is employed. This representation facilitates the evaluation of the rover's motion relative to its own frame, offering a clearer understanding of its alignment with the planned path. By combining absolute and body-centric poses, the navigation system effectively reconstructs the rover's trajectory while ensuring an intuitive interpretation of its movement within the simulation. This dual representation enhances localization accuracy and provides a structured framework for analyzing positioning errors.

Although the present formulation is planar, the proposed architecture is inherently compatible with an extension to full 6-DoF motion. In that case, the rover pose would be expressed as $(x, y, z, \phi, \theta, \psi)$, including both translational and rotational components. Such an extension would require preserving the full 3D structure of the reference frames, propagating roll and pitch within the inertial navigation stage, and extending the fusion process so that IMU, MC, and LiDAR provide complete spatial pose estimates rather than planar projections. Error analysis would also be extended to include full 3D position and orientation components.

3 Results and Discussion

In this section, the simulation setup is detailed and the results achieved are evaluated. The discussion begins with an overview of the simulation tools and their specific roles. The analysis then follows the structure outlined in Fig. 2, illustrating how each module is implemented and examining the corresponding outputs in the order they were introduced. A comparative analysis with a standard odometry-based EKF is then provided, followed by an exploration of the EKF's sensitivity to variations in the Q and R matrices.

3.1 Simulation Framework

The simulation framework includes Blender™, Unreal Engine™, MATLAB™, and Simulink™.

Blender™ is used to refine the rover's 3D model and perform camera calibration using a checkerboard pattern. The selected MC is the FLIR™ Blackfly S¹, featuring a 90° horizontal field of view (hFOV) and 1.3 MPx resolution. Calibration is carried out in MATLAB™ using synthetic images from Blender™ to estimate intrinsic and extrinsic parameters.

Unreal Engine™ provides a high-fidelity simulation environment and handles real-time interaction between the rover and the virtual world. The setting is kept simple and flat to allow controlled testing and to ensure consistency with the planar-motion assumption adopted in the navigation framework. The imported rover model is configured with blueprints for the MC and LiDAR actors, which interact with Simulink™ via predefined tags. The IMU, on the other hand, is not simulated within Unreal Engine™. Instead, it is modeled in Simulink™ to enable a dedicated acquisition rate, independent of the rendering frequency, thus avoiding the inaccuracies associated with frame-based sampling.

Simulink™ handles rover dynamics and sensor data acquisition, acting as the interface between Unreal Engine™ and MATLAB™ via the Vehicle Dynamics Blockset [25] and the Automated Driving Toolbox. A global simulation model organizes all components, with specific blocks for sensor transforms and scene setup. Sensor positions are defined relative to the rover, with LiDAR placed above the MC to maximize field of view (FOV) and reduce interference. The final vertical field of view (vFOV) is limited to 30° through an outlier removal filter.

MATLAB™ acts as the computational core, managing SLAM execution, trajectory generation, and data processing. Its integration with the other platforms ensures coherence across the full pipeline.

3.2 Simulation Part I: Data Harvesting

Data generation is divided into simulation setup and sensor data collection. The first phase defines waypoints on a 2D scene, smooths the trajectory, and generates velocity and heading inputs using the odometry motion model, previously detailed in Section 2.2.1, which decomposes movement into initial rotation, translation, and final rotation. To ensure smooth heading transitions, yaw angles in the $[0^\circ, 360^\circ]$ range are unwrapped in order to avoid discontinuities. During data collection, LiDAR point clouds, MC image frames, and IMU acceleration and angular velocity are stored in separate files. These raw data will be used in the next phase for trajectory reconstruction and fusion.

3.3 Simulation Part II: First Layer of the Navigation Filter - Data Processing

In the first filtering layer, raw data from IMU, MC, and LiDAR are processed independently to obtain pose estimates. Each sensor receives raw data from the DIM and computes its own estimate without cross-referencing others.

¹<https://www.flir.it/products/blackfly-s-usb3/?model=BFS-U3-13Y3C-C&vertical=machine+vision&segment=iis>

The IMU data, composed of accelerometer and gyroscope measurements, are used to reconstruct the rover’s trajectory. Angular velocity is integrated to estimate yaw, which is then synchronized with MC timestamps. Linear acceleration is transformed into the ARF using the estimated heading, and integrated to retrieve x and y positions. The inputs to this process are the raw IMU measurements (linear acceleration and angular velocity) acquired in the rover’s body frame. The output is a pose trajectory in the form (x, y, θ) referred to the ARF, obtained through integration without the aid of external references.

For the MC, a modified Oriented FAST and Rotated BRIEF SLAM (ORB-SLAM) extracts keyframes and builds a sparse 3D map. ORB features are matched across frames, and the resulting poses are refined via bundle adjustment. The estimated trajectory is scaled, rotated, and translated to align with ground truth and corrected for the MC offset to obtain rover poses in ARF. In this case, the input consists of the camera image frames and intrinsic parameters. The output is a corrected sequence of rover poses in the ARF, expressed as (x, y, θ) , derived from the estimated MC trajectory and improved through extrinsic calibration.

Regarding the LiDAR, point clouds are preprocessed by removing ground planes and outliers. Key LiDAR Odometry and Mapping (LOAM) features (edges and planes) are extracted and downsampled. A LOAM-based odometry system estimates the LiDAR pose. Since the LiDAR is rigidly attached to the rover, its yaw is assumed to coincide with the rover yaw. Map points and trajectory data are retrieved in the ARF. The input is the raw point cloud stream from the sensor. The output includes both the estimated rover trajectory in the ARF (x, y, θ) together with a preliminary 3D map of the environment, that will be used in later stages for mapping refinement.

The next paragraphs evaluate each sensor’s performance in terms of accuracy, noise, and completeness of the reconstructed trajectory.

3.3.1 IMU Data Processing Results

The IMU, operating solely within the rover’s body coordinates, provides direct motion data without reference to the external environment. To emulate realistic sensor behavior, different types of noise are introduced into the IMU measurements. These include white noise, representing random measurement error, and a constant bias applied to each axis as detailed in Tab. 1. Drift naturally arises from the integration of biased and noisy inertial measurements over time.

Table 1 IMU noise parameters

Accelerometer	
Constant bias (m/s ²)	(0.08, -0.05, 0.03)
Noise power ((m/s ²) ² /Hz)	(0.0015, 0.0015, 0.0015)
Gyroscope	
Constant bias (rad/s)	(0.002, -0.002, 0.001)
Noise power ((rad/s) ² /Hz)	(0.00015, 0.00015, 0.00015)

The resulting trajectory and heading angle estimations highlight the limitations of using IMU alone, reinforcing the need for sensor fusion due to visible drift over time. This behavior is illustrated in Fig. 3, showing the reconstructed poses deviating progressively from the ground truth.

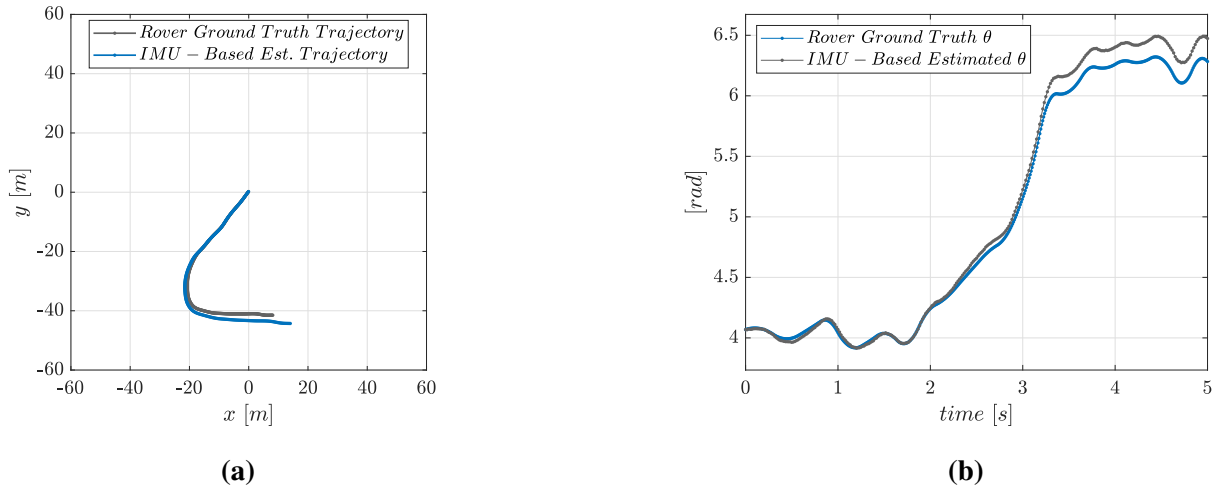


Fig. 3 Final rover poses reconstruction using IMU data: (a) Trajectory, (b) Heading

3.3.2 MC Data Processing Results

The MC data are initially processed through the ORB-SLAM pipeline, producing a trajectory in the sensor’s reference frame, lacking absolute scale. A transformation aligns it with the WCS and later with the ARF, accounting for the camera’s mounting offset. The results show limited drift and reasonable accuracy in both position and heading, confirming the MC’s reliability when supported by proper calibration and post-processing. The estimated pose evolution, compared against the ground truth, is displayed in Fig. 4.

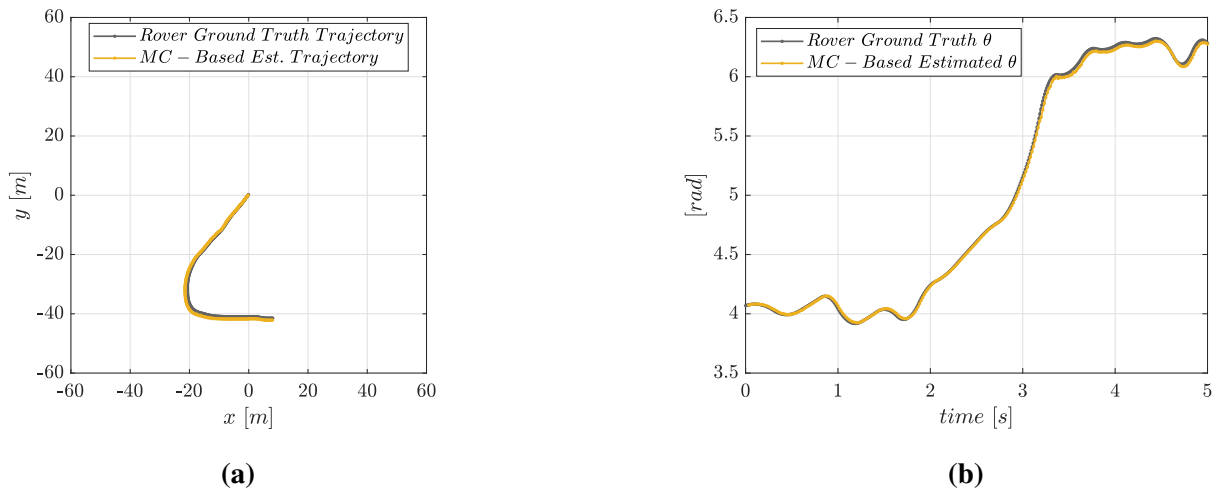


Fig. 4 Final rover poses reconstruction using MC data: (a) Trajectory, (b) Heading

3.3.3 LiDAR Data Processing Results

The LiDAR proves to be the most accurate sensor in both localization and mapping. Following preprocessing steps such as ground plane removal and segmentation, the LOAM pipeline reconstructs the rover’s trajectory and a dense environmental map. After correcting for mounting offsets, the LiDAR-based trajectory is transformed into the ARF and exhibits minimal deviation from ground truth, making it an essential component of the SLAM pipeline. The final map, after reference alignment, offers a highly consistent spatial reconstruction. Fig. 5 confirms that the reconstructed poses closely follow the simulated reference trajectory.

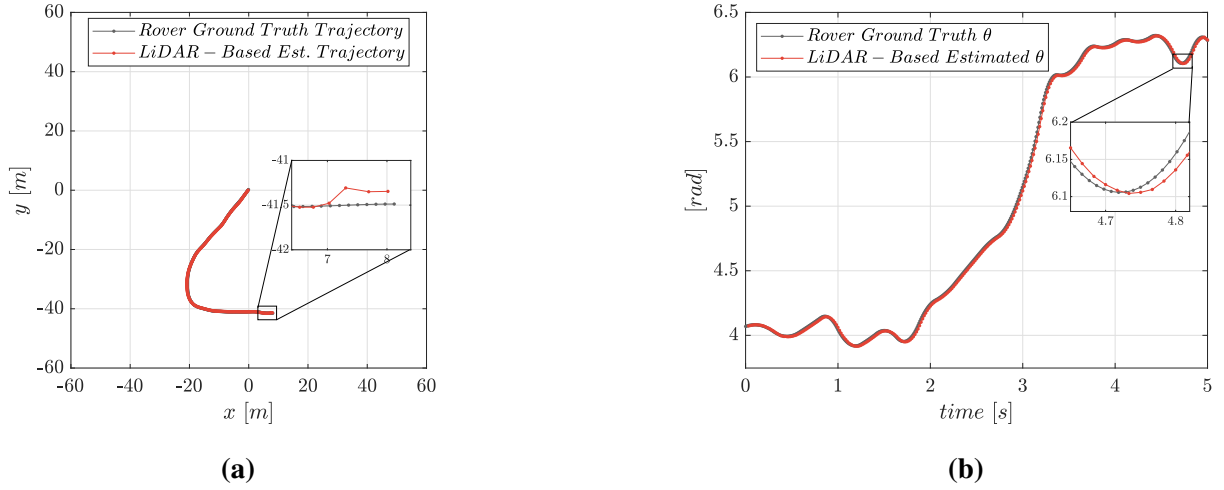


Fig. 5 Final rover poses reconstruction using LiDAR data: (a) Trajectory, (b) Heading

3.4 Simulation Part III: Second Layer of the Navigation Filter - Sensor Fusion

This stage fuses the pose estimates produced by individual sensors into a single, optimized trajectory through a two-step process.

The first loop employs a modified GB method that, unlike traditional approaches such as GraphSLAM, does not incorporate landmarks. Rather, it relies on independent trajectory estimates from the IMU and the MC, which are fused into a unified representation. This independence from mapped features and explicit control inputs makes the approach particularly advantageous for scenarios where landmark-based mapping is impractical or unreliable. The fusion strategy is implemented locally at each time step by combining the independent pose estimates from the two sensors through a weighted averaging process. Each sensor is associated with a confidence matrix that determines its contribution to the fused estimate. These weights are selected to reflect the relative reliability of the IMU and MC measurements, based on their observed behavior in the simulation environment rather than on a formal probabilistic model. In particular, the IMU provides accurate short-term motion estimation but suffers from drift over time, while the MC exhibits more stable long-term behavior despite being affected by systematic errors. The assigned weights therefore balance these complementary characteristics, with the camera contributing to long-term consistency and the IMU supporting local trajectory continuity. The confidence matrices are implemented as diagonal weighting matrices applied to each component of the pose vector, allowing different weights to be assigned to position and heading. Fusion is then performed locally at each time step through a weighted averaging of the independent pose estimates. Specifically, if both estimates are available, the optimized pose $\hat{\mathbf{p}}$ is computed as:

$$\hat{\mathbf{p}} = (\mathbf{W}_{\text{IMU}} + \mathbf{W}_{\text{MC}})^{-1} (\mathbf{W}_{\text{IMU}} \mathbf{p}_{\text{IMU}} + \mathbf{W}_{\text{MC}} \mathbf{p}_{\text{MC}}) \quad (2)$$

where \mathbf{p}_{IMU} and \mathbf{p}_{MC} are the pose estimates from each sensor, and \mathbf{W}_{IMU} , \mathbf{W}_{MC} are the associated confidence matrices. When only a single estimate from one of the two sensors is available, its corresponding output is used to ensure continuity. This weighted averaging can be interpreted as the closed-form solution of a local least-squares problem, where the fused pose minimizes the residuals with respect to the individual sensor estimates. The associated cost function can be written as:

$$J(\mathbf{p}) = (\mathbf{p} - \mathbf{p}_{\text{IMU}})^T \mathbf{W}_{\text{IMU}} (\mathbf{p} - \mathbf{p}_{\text{IMU}}) + (\mathbf{p} - \mathbf{p}_{\text{MC}})^T \mathbf{W}_{\text{MC}} (\mathbf{p} - \mathbf{p}_{\text{MC}}) \quad (3)$$

The analytical solution is obtained by minimizing $J(\mathbf{p})$ with respect to \mathbf{p} , i.e., setting its gradient to zero, which yields the expression reported in Eq. (2).

The method can be classified as graph-based due to its structure and update strategy. The trajectory is modeled as a graph, where each node represents a pose estimate, and edges encode constraints derived from the IMU and the MC. These constraints define the relative motion between consecutive poses and are incorporated into a weighted averaging process. This formulation ensures that the optimized pose remains closer to the more reliable sensor's estimate while still incorporating the complementary data from the other. Mathematically, this process minimizes the weighted sum of residuals, i.e., the difference between the optimized pose and each sensor estimate. Although the optimization occurs locally for each pose rather than globally across the entire trajectory, the method adheres to the graph-based principle of enforcing constraints between connected nodes. The resulting trajectory reflects the strengths of both sources and serves as input for the EKF refinement loop.

The second loop further improves the optimized trajectory from the GB method using an EKF. The latter operates on a reduced state vector $\mathbf{p} = (x, y, \theta)$ and uses LiDAR measurements to correct the GB estimates. The EKF initialization phase involves defining the state covariance matrix P , the process noise covariance Q , and the measurement noise covariance R . The LiDAR pose estimates are temporally aligned with the GB trajectory to ensure consistent data fusion at each time step. When no LiDAR measurement is available, the state is propagated using the GB trajectory, and the uncertainty is increased to reflect model assumptions by adding Q to the current P . This behavior is consistent with the overall structure of the fusion framework, which is designed to preserve continuity even when one of the sensor estimates is temporarily unavailable. In the first loop, this is achieved by using the available IMU or MC estimate whenever the other is missing, while in the second loop the EKF propagates the state using the GB trajectory whenever no LiDAR update is provided. In the current setup, LiDAR measurements are available for almost the entire trajectory, except for the initial discarded acquisitions. For this reason, prolonged LiDAR outage conditions were not explicitly simulated, although the formulation remains compatible with temporary sensor unavailability. Once LiDAR data becomes available, a complete prediction-update cycle is performed. The predicted state and its covariance are initialized as:

$$\mathbf{p}_{\text{pred}} = \hat{\mathbf{p}}, \quad P_{\text{pred}} = P + Q \quad (4)$$

The update step integrates the LiDAR measurement \mathbf{z} . The observation model is defined as a direct measurement of the rover pose in the ARF:

$$\mathbf{z} = H\mathbf{p} + \mathbf{v}, \quad H = I \quad (5)$$

where \mathbf{v} represents zero-mean measurement noise with covariance R . This reflects the fact that the LiDAR processing pipeline provides pose measurements in the same state space (x, y, θ) , thus enabling a direct correction of the state without the need for a nonlinear observation mapping. The Kalman gain is computed as:

$$K = P_{\text{pred}}H^T (HP_{\text{pred}}H^T + R)^{-1} \quad (6)$$

The state is then updated as:

$$\mathbf{p}_{\text{est}} = \mathbf{p}_{\text{pred}} + K(\mathbf{z} - H\mathbf{p}_{\text{pred}}) \quad (7)$$

and the posterior covariance becomes:

$$P = (I - KH)P_{\text{pred}} \quad (8)$$

This formulation corresponds to a linear observation model, where the measurement constrains the state estimate, leading to a correction step that minimizes the innovation weighted by the measurement uncertainty. Unlike traditional EKF-based SLAM, which jointly estimates both poses and landmarks, this implementation simplifies the state, excluding explicit landmark tracking. Instead, LiDAR measurements are treated as pose observations in the update step, reducing reliance on a control model. This approach leverages the high precision of LiDAR, which consistently outperforms IMU and MC in measurement

accuracy. The prediction step uses the GB trajectory to propagate the rover’s state forward, while the update step refines this estimate using LiDAR measurements through the Kalman gain. This reduces dependency on odometry-based motion models, which are often prone to drift and require control inputs such as wheel velocities.

A final error analysis, regarding both loops, is conducted in the RRF, computing deviations in along-track s , cross-track n , and heading θ , with corresponding 3σ confidence bounds. These metrics summarize the localization accuracy achieved after both fusion steps, whose effectiveness is further evaluated in the following sections.

3.4.1 Graph-Based Method Results

The GB method fuses IMU and MC pose estimates, balancing their contributions through confidence-based weighted averaging. The resulting trajectory aligns closely with the MC estimate while reducing cumulative drift thanks to the IMU’s high-frequency updates. As illustrated in Fig. 6, the positional and heading errors in the RRF are significantly lower than those from individual sensors, confirming the method’s effectiveness.




Fig. 6 Comparison between GB error and cumulative sensor errors in RRF

3.4.2 Extended Kalman Filter Results

The EKF further refines the GB trajectory by integrating LiDAR pose estimates. It operates on a reduced state and does not require control inputs. The choice of the EKF’s noise covariance matrices was guided by the sensitivity analysis discussed later in Section 3.6, aiming to balance prediction responsiveness and measurement reliability. In particular, the measurement noise configuration reflects the expected performance of the LiDAR sensor, while the process noise ensures adaptability without overfitting. Fig. 7 illustrates the EKF’s final trajectory and heading estimation. These results confirm the filter’s effectiveness in smoothing the pose reconstruction and maintaining temporal consistency in a GPS-denied setting. The accuracy achieved through this refinement process will be further examined in the following section.




Fig. 7 EKF results: (a) Trajectory estimation, (b) Heading estimation

3.5 Comparative Analysis

To evaluate the benefits of replacing control-based prediction with the GB trajectory, a comparison is made between the proposed method (EKF-GB) and a standard EKF implementation using odometry inputs (EKF-Odometry). Both filters share the same initialization parameters, ensuring that the only difference lies in the prediction model.

In EKF-Odometry, the prediction step uses noisy control inputs, which alone would result in a significant deviation from the true trajectory. Although LiDAR updates partially correct the trajectory, cumulative errors in heading and lateral displacement often remain. By contrast, the EKF-GB method leverages a fused IMU-MC trajectory as prior, which provides a more stable and accurate basis for prediction. Tab. 2 summarizes the average errors obtained for both configurations. The EKF-GB approach consistently outperforms its odometry-based counterpart across all components of the pose, with reductions of approximately 4% in along-track error, over 25% in cross-track error, and more than 20% in heading error. These results highlight the value of adding a lightweight graph-based layer before the EKF, especially when odometry is unreliable. The reduced errors suggest improved trajectory consistency and better alignment with the ground truth path, making the EKF-GB approach particularly suitable for long-term navigation in unstructured or GPS-denied environments.

Table 2 Mean errors comparison

Mean error	EKF-GB	EKF-Odometry
$\overline{\text{err}}_s$ (m)	0.2920	0.3041
$\overline{\text{err}}_n$ (m)	0.0414	0.0558
$\overline{\text{err}}_\theta$ (deg)	0.6245	0.8005

While the observed gains might appear moderate, it is important to consider the context. Due to the high computational load imposed by the simulation framework, especially the integration of Unreal Engine™ and the detailed sensor modeling, tests were necessarily limited to short time spans. Nonetheless, the EKF-GB method demonstrated measurable improvements. It is reasonable to expect that, over longer mission durations, the benefits of the proposed approach would become even more pronounced.

3.6 EKF Dependence on Noise Covariance Matrices

The impact of Q and R was also investigated. Lower values of Q reduce corrections from LiDAR, making the EKF overly reliant on prediction, while higher values allow more effective integration of measurements. Conversely, low R values overemphasize sensor data, which can be detrimental if measurements are noisy, while high R values reduce the filter's responsiveness.

The evolution of the 3σ uncertainty bounds shows that a balanced setting leads to rapid convergence and long-term consistency. Setting $Q = 0$ slows convergence and may cause the filter to exceed expected error thresholds. These findings suggest the potential of an adaptive tuning strategy for R based on residuals analysis. Such a solution could improve robustness and adaptability, enhancing performance.

4 Conclusions

This study presented a modular SLAM framework for lunar UGVs operating in GPS-denied environments, developed within a flexible simulation architecture integrating multiple software platforms and heterogeneous sensors. This setup enabled realistic development and validation of SLAM algorithms for lunar exploration.

Results demonstrate the effectiveness of the approach, with accurate trajectory reconstruction and error metrics consistently within 3σ bounds. Compared to standard EKF odometry methods, the proposed formulation achieved superior accuracy and stability. The system's independence from control inputs, odometry models, and global landmarks enhances reliability against drift and feature scarcity. Combined with its modularity, the EKF-GB approach offers a versatile, resilient alternative adaptable to evolving mission needs.

Future work includes adaptive tuning of noise covariances, exploring alternative filters such as the Unscented Kalman Filter, enhancing simulation realism through varied terrains and dynamic obstacles, incorporating additional sensors (e.g., radar and thermal cameras), extending the current planar formulation to full 6-DoF pose estimation, and broadening the architecture toward multi-robot systems for improved mapping and operational resilience. In particular, the extension to 6-DoF motion will require modelling non-planar rover dynamics, including vertical displacement, roll, and pitch, thereby enabling this approach to address more realistic lunar surface and subsurface scenarios characterized by slopes, irregular ground contact, and body attitude variations.

These directions will further strengthen the method as a robust foundation for autonomous navigation in complex, unstructured environments.

Acknowledgments

This work is largely inspired by the author's Master's Thesis [26] and further developed within the research framework established at Politecnico di Milano.

The author wishes to express sincere gratitude to Professor Mauro Massari, supervisor at Politecnico di Milano, for his invaluable guidance and constant support, providing the knowledge to tackle complex problems and fostering skills that proved essential in the development of this work. Special thanks are also extended to co-supervisor Annachiara Ippolito, whose knowledge, availability, and encouragement greatly contributed to the progress of this research.

Declaration of Use of Artificial Intelligence

Artificial intelligence was used exclusively to assist with proofreading, translation, and linguistic reorganization of the manuscript. The authors wrote the original text and provided the complete scientific content, including research design, data analysis, coding, and manuscript structure. The use of AI did not contribute to the generation of scientific results, images, or any substantive content. The authors remain fully responsible for the integrity of the work.

References

- [1] Ronald Greeley. Lava tubes and channels in the lunar Marius Hills. *The Moon*, 3(3):289–314, Dec. 1971. doi: [10.1007/BF00561842](https://doi.org/10.1007/BF00561842).
- [2] Raymond P. Martin and Haym Benaroya. Pressurized lunar lava tubes for habitation. *Acta Astronautica*, 204:157–174, 2023. doi: [10.1016/j.actaastro.2022.12.013](https://doi.org/10.1016/j.actaastro.2022.12.013).
- [3] Xiaohang Qiu and Chunyu Ding. Radar observation of the lava tubes on the moon and mars. *Remote Sensing*, 15(11), 2023. doi: [10.3390/rs15112850](https://doi.org/10.3390/rs15112850).
- [4] Audai K. Theinat, Anahita Modiriasari, Antonio Bobet, H. Jay Melosh, Shirley J. Dyke, Julio Ramirez, Amin Maghareh, and Daniel Gomez. Lunar lava tubes: Morphology to structural stability. *Icarus*, 338:113442, 2020. doi: [10.1016/j.icarus.2019.113442](https://doi.org/10.1016/j.icarus.2019.113442).
- [5] Junichi Haruyama, Kazuyuki Hioki, Motomaro Shirao, Tomokatsu Morota, Harald Hiesinger, Carolyn H. van der Bogert, Hideaki Miyamoto, Akira Iwasaki, Yasuhiro Yokota, Makiko Ohtake, Tsuneo Matsunaga, Seiichi Hara, Shunsuke Nakanotani, and Carle M. Pieters. Possible lunar lava tube skylight observed by SELENE cameras. *Geophysical Research Letters*, 36(21), 2009. doi: [10.1029/2009GL040635](https://doi.org/10.1029/2009GL040635).
- [6] William Whittaker. Technologies enabling exploration of skylights, lava tubes and caves. In *NASA Innovative Advanced Concepts (NIAC) Phase*, 2012. <https://api.semanticscholar.org/CorpusID:52206458>.
- [7] Filip Klaesson, Petter Nilsson, Tiago Stegun Vaquero, Scott Tepsuporn, Aaron Ames, and Richard M. Murray. Planning and optimization for multi-robot planetary cave exploration under intermittent connectivity constraints. In *Proceedings (venue unspecified)*, 2020. <https://api.semanticscholar.org/CorpusID:231797539>.
- [8] Jet Propulsion Laboratory and NASA Planetary Science Division. Guidance, navigation, and control technology assessment for future planetary science missions, part iii: Surface and subsurface guidance, navigation, and control. Technical Report JPL D-110031, NASA / Jet Propulsion Laboratory, 2023.
- [9] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA, 2005. ISBN: 978-0-262-20162-9.
- [10] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localisation and mapping (slam): Part i the essential algorithms. *IEEE Robotics & Automation Magazine*, 13(1):99–110, Jan. 2006. doi: [10.1109/MRA.2006.1638022](https://doi.org/10.1109/MRA.2006.1638022).
- [11] Cyrill Stachniss, John J. Leonard, and Sebastian Thrun. Simultaneous localization and mapping. In Bruno Siciliano and Oussama Khatib, editors, *Springer Handbook of Robotics*, pages 1153–1176. Springer, Cham, Switzerland, 2016. doi: [10.1007/978-3-319-32552-1_46](https://doi.org/10.1007/978-3-319-32552-1_46).
- [12] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, Jose Neira, Ian Reid, and John J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, Dec. 2016. doi: [10.1109/TRO.2016.2624754](https://doi.org/10.1109/TRO.2016.2624754).
- [13] Randall Smith, Matthew Self, and Peter Cheeseman. Estimating uncertain spatial relationships in robotics. In Ingemar J. Cox and Gordon T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer, New York, NY, USA, 1990. doi: [10.1007/978-1-4613-8997-2_14](https://doi.org/10.1007/978-1-4613-8997-2_14).



- [14] Rathin Chandra Shit. Precise localization for achieving next-generation autonomous navigation: State-of-the-art, taxonomy and future prospects. *Computer Communications*, 160:351–374, 2020. doi: [10.1016/j.comcom.2020.06.007](https://doi.org/10.1016/j.comcom.2020.06.007).
- [15] Jeffrey Martz, Wesam Al-Sabban, and Ryan N. Smith. Survey of unmanned subterranean exploration, navigation, and localisation. *IET Cyber-Syst. Robot.*, 2(1):1–13, 2020. doi: [10.1049/iet-csr.2019.0043](https://doi.org/10.1049/iet-csr.2019.0043).
- [16] Jun Zhu, Hongyi Li, and Tao Zhang. Camera, LiDAR, and IMU based multi-sensor fusion SLAM: A survey. *Tsinghua Science and Technology*, 29(2):415–429, 2024. doi: [10.26599/TST.2023.9010010](https://doi.org/10.26599/TST.2023.9010010).
- [17] Héctor Azpúrua, Maíra Saboia, Gustavo M. Freitas, Lillian Clark, Ali-akbar Agha-mohammadi, Gustavo Pessin, Mario F. M. Campos, and Douglas G. Macharet. A survey on the autonomous exploration of confined subterranean spaces: Perspectives from real-world and industrial robotic deployments. *Robotics and Autonomous Systems*, 160:104304, 2023. doi: [10.1016/j.robot.2022.104304](https://doi.org/10.1016/j.robot.2022.104304).
- [18] Jet Propulsion Laboratory and NASA Planetary Science Division. Guidance, navigation, and control technology assessment for future planetary science missions, part iv: Aerial guidance, navigation, and control. Technical Report JPL D-110047, NASA / Jet Propulsion Laboratory, 2023.
- [19] European Space Agency. Esa plans mission to explore lunar caves, 2021. Accessed: 2025. https://www.esa.int/Enabling_Support/Preparing_for_the_Future/Discovery_and_Preparation/ESA_plans_mission_to_explore_lunar_caves.
- [20] Chris Dinelli, John Racette, Mario Escarcega, Simon Lotero, Jeffrey Gordon, James Montoya, Chase Dunaway, Vasileios Androulakis, Hassan Khaniani, Sihua Shao, Pedram Roghanchi, and Mostafa Hassanalian. Configurations and applications of multi-agent hybrid drone/unmanned ground vehicle for underground environments: A review. *Drones*, 7(2), 2023. doi: [10.3390/drones7020136](https://doi.org/10.3390/drones7020136).
- [21] J. Balaram, T. Canham, C. Duncan, M. Golombek, H. F. Grip, W. Johnson, J. Maki, A. Quon, R. Stern, and D. Zhu. Mars helicopter technology demonstrator. In *AIAA Atmospheric Flight Mechanics Conference*, 2018.
- [22] Ryan Soussan, John McCaffery, Scott McMichael, and Matthew Deans. Luvo: Lunar visual odometry using homography-based image feature matching. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 13428–13435, 2025. doi: [10.1109/ICRA55743.2025.11127410](https://doi.org/10.1109/ICRA55743.2025.11127410).
- [23] Weifeng Chen, Chengjun Zhou, Guangtao Shang, Xiyang Wang, Zhenxiong Li, Chonghui Xu, and Kai Hu. Slam overview: From single sensor to heterogeneous fusion. *Remote Sensing*, 14(23), 2022. doi: [10.3390/rs14236033](https://doi.org/10.3390/rs14236033).
- [24] Automated driving toolbox, 2025. Online; accessed 2025. <https://it.mathworks.com/products/automated-driving.html>.
- [25] Vehicle dynamics blockset, 2025. Online; accessed 2025. <https://it.mathworks.com/products/vehicle-dynamics.html>.
- [26] Niccolò Goretti. Enhanced sensor fusion for SLAM-driven autonomous navigation of lunar UGVs in GPS-denied environments. Master’s thesis, Politecnico di Milano, Dept. of Aerospace Science and Technology, Milan, Italy, 2024.