

A Dual-Mode Approach for Vision-Based Navigation in a Lunar Landing Scenario

Luca Ostrogovich^{1,2} Roberto Del Prete^{1*} Giuseppe Tomasicchio²
Nicolas Longépé³ Alfredo Renga¹

¹Department of Industrial Engineering, University of Naples Federico II

²Telespazio SRL

³Φ-lab, ESA ESRIN

{luca.ostrogovich; roberto.delprete; alfredo.renga}@unina.it

giuseppe.tomasicchio@telespazio.com nicolas.longepe@esa.int

Abstract

In this research, a novel approach for autonomous spacecraft navigation, particularly in lunar contexts, is presented, focusing on vision-based techniques. The system incorporates lunar crater recognition in conjunction with feature tracking to enhance the accuracy of spacecraft navigation. This system underwent comprehensive evaluation through a purpose-built software simulation, replicating lunar conditions for thorough evaluation and refinement. The methodology integrates established navigational methods with advanced artificial intelligence algorithms, resulting in significant navigational accuracy. The system demonstrates precise capabilities in determining the spacecraft position, with an average accuracy of approximately 270 m for the absolute navigation mode, while the relative mode exhibited an average error of 27.4 m and 0.8 m in determining the horizontal and vertical lander displacements relative to terrain. Initial tests on embedded systems—akin to those on-board spacecraft—were conducted. These tests are pivotal in demonstrating the system’s operational viability within the constraints of limited bandwidth and rapid processing requirements characteristic of space missions. The promising results from these tests suggest potential applicability in real-world space missions, enhancing autonomous navigation capabilities in lunar and potentially other extraterrestrial environments.

1. Introduction and Background

The Lunar exploration venture has gathered renowned attention in recent years. The scientific interest in lunar exploration, especially in the polar regions with confirmed ice, is spurring efforts to establish bases and conduct sampling [3]. On February 22 2024, the Nova-C lander, named

Odysseus and developed by Intuitive Machines, successfully concluded a seven-day orbital transit to the Moon. It achieved a soft landing in proximity to Malapert A crater, located within the Moon’s South Pole region, at 18:24 Eastern Standard Time¹. The revival in manned moon missions challenges navigation systems, as space missions face high costs and operational limits with networks like DSN or ESTRACK. Additionally, unreliable satellite links and inaccessible ground stations complicate real-time command transmission. Notwithstanding, acquiring knowledge of the spacecraft’s position and orientation is imperative for navigating around a planetary body or conducting In-Orbit Servicing (IOS) and Active Debris Removal (ADR) operations. In these challenging scenarios, optical navigation (OPNAV) methods can be a viable tool for determining the spatial orientation and coordinates of a spacecraft, as well-documented in the literature [16]. Indeed, Visual-Based Navigation (VBN) strategies are implemented to ascertain a spacecraft’s spatial orientation and coordinates. This methodology leverages image data and is bifurcated into two primary categories, contingent upon the relative proximity of the spacecraft to the celestial body [2]. Firstly, the horizon-based approach is pertinent when a substantial segment of the celestial body is discernible within the captured imagery. Conversely, as the spacecraft approaches the vicinity of the target surface, the focus shifts to landmark-based navigation [7]. This technique involves the identification and utilization of distinct topographical features, such as craters, boulders, or other unique terrain elements, for precise positional determination [2]. In the ambit of VBN, various sensor options exist, including monocular stereo cameras, thermal sensors RGB/greyscale cameras, RADARs, LIDARs, among others. Monocular cameras are

¹<https://www.nasa.gov/missions/lro/nasas-lro-images-intuitive-machines-odysseus-lander/>

often preferred over other active sensors such as LIDARs and RADARs due to their relative simplicity, compact size, lightweight nature, minimal power requirements, and ease of integration into a wide array of spacecraft configurations. Over the past decade, spacecraft navigation based on vision has employed hand-crafted features described using feature descriptors and identified through feature detectors. The recent leaps in Deep Learning (DL) have led to substantial advancements in the development of computer vision algorithms, enhancing their efficacy and robustness across various applications, including image classification [41] and segmentation [26]. Reflecting this trend, the rise of DL-based spacecraft pose estimation algorithms has recently surpassed that of traditional feature-engineering-based methods [29]. Indeed, the application of Artificial Intelligence (AI) in VBN addresses several challenges, such as noise, variable lighting, and changing shadow patterns in imaging systems. AI's adaptability to these conditions enhances intelligent data analysis, and interpretation from onboard sensors, aiding rapid and well-informed decision-making. As stated above, locating a spacecraft flying near a celestial body with a monocular camera usually involves employing object detection techniques tailored to identifying surface landmarks. DL-based Object detectors play a crucial role in this effort, facilitating the detection of specific features on the celestial body's surface, and have been widely implemented in lunar crater detection tasks for navigation and geological purposes [40] [33].

The AI integration in spacecraft brings distinct challenges, particularly the limitation of low power availability, which impacts the choice of AI models. Effective real-time performance of an AI model requires meticulous selection of its architecture. This involves strategies to reduce inference times and computational complexity, including model quantization and distillation. The development of AI-specific processors is helping to overcome these barriers, enabling the practical use of AI in both terrestrial and onboard satellite applications. Nowadays, AI has been integrated into various Earth Observation (EO) missions. The ϕ -Sat-1 [15], equipped with an AI-accelerated chip for cloud detection, pioneered this integration. This development was followed by other significant missions like HYPPO-1 [1], WildRide [14, 25], and OPS-SAT [17], each serving different purposes.

In the framework of VBN for lunar landing, numerous solutions employing AI methodologies have been proposed. Nonetheless, their practical application onboard satellites has been evaluated in rare cases. For instance, the strategy employed by Downes et al. [8], reformulated the crater detection challenge into a segmentation task, utilizing a U-net architecture. Their U-net, based on the architecture of DeepMoon [33], which detected lunar craters from elevation imagery, underwent further training with 800 Lu-

nar Reconnaissance Orbiter Camera (LROC) intensity images. The identified craters are treated as features that are tracked using an Extended Kalman Filter (EKF). However, this segmentation-based method tends to be inefficient in terms of computational resource utilization. The inherent increase in parameters due to the decoder component escalates the inference time, a factor critical to consider for onboard systems. Moreover, a segmentation approach restricts the integration of alternative image processing techniques such as dilation, erosion, and ellipse fitting for crater detection. These additional steps further increase latency and become superfluous when utilizing an optimally designed network architecture capable of directly addressing the detection task.

In [34], a crater detector employing the Single Shot Detector (SSD) [23] architecture coupled with MobileNetV2 [31] has been utilised to highlight craters in images of the lunar surface. This detector was trained via transfer learning, leveraging a simulated portion of the Moon derived from a Digital Elevation Model (DEM) generated in Pangu [24]. The network has been implemented on a Raspberry Pi 4 equipped with an Intel® Movidius™ Neural Compute Stick, achieving an inference time of 0.4 seconds per image. To the best of the authors' knowledge, the work discussed herein represents the closest parallel to our own. Notwithstanding, there are notable distinctions: a) the referenced study utilized simulated imagery rather than authentic Lunar Reconnaissance Orbiter (LRO) images for training the crater detector; b) the synthetic images are acquired over a handcrafted lunar surface portion, and not over a high-fidelity surface reconstructed from the LRO Laser Altimeter (LOLA) generated DEM; and c) it employed a singular mode (absolute) which may not be optimally suited for lower altitudes where crater prevalence is limited. In this manuscript, we advance the field by introducing a dual-mode system designed to effectively address the scarcity of craters. In summary, the principal contributions of this research are as follows:

1. **Development of a novel dual-mode VBN system:** we propose a unique VBN system with dual operational modes, tailored to the spacecraft's varying altitude.
2. **System evaluation:** we conduct an in-depth analysis of the proposed system, focusing on accuracy and timing in position estimation within a targeted trajectory.
3. **Model testing on embedded hardware:** Focusing on the AI component, this research includes preliminary testing of the crater detection module on specialized hardware, specifically evaluating the algorithm's efficiency and latency.

The structure of the following sections of the manuscript is arranged as follows: Section 2 presents an exhaustive delineation of the proposed system architecture, elaborating on each constituent component in detail. Section 3 elucidates

the designated target trajectory and imparts numerical findings on navigation performance across the examined modalities. Section 4 furnishes a comprehensive synopsis of the current progress and delineates the requisite steps for the system’s integration onboard. The manuscript culminates with Section 5, which provides a conclusive summary.

2. Dual-Mode VBN

In the lunar environment, craters are the perfect landmark candidates due to their distinct shape and morphology, which results in their rims casting sharp shadows. Craters represent a particular type of visual landmark present in the lunar environment that can be observed under various lighting conditions and with a diverse array of sensors. Presently, crater databases encompass the majority of craters with diameters larger than 1–2 km [30]. It is therefore easy to envision a navigation system that leverages these topographical markers and correlates them with an extensive catalogue to achieve terrain-based navigation. Many DL-based lunar crater detectors have been proposed [8] [33] [34] [40].

As the lander’s altitude decreases, however, fewer and fewer craters are framed by the camera; moreover, the craters that are framed are too small to be surveyed in a lunar crater catalogue. It follows that this approach well suits the preliminary descent phase, from the 100 km Low Lunar Orbit (LLO) to an altitude of about 10 km, where many, distinct and catalogued craters are visible in the sensor’s FOV. Hence, as the lander’s altitude descends below 10 km, a shift to a relative visual navigation algorithm becomes necessary. This relative operational mode would no longer rely on crater-based markers but instead on tracking common surface features, like rocks and fissures, across subsequent frames. The proposed visual navigation pipeline (Figure 1) will provide navigation support for the entire descent trajectory, including segments where a crater matching strategy is no longer possible. The following subsections provide comprehensive details on the two distinct operational modes.

2.1. Absolute Visual-based Navigation

The proposed crater-based absolute navigation algorithm unfolds across three main stages. Initially, high-resolution onboard imagery from a spacecraft is utilized as input for a Deep Learning-based crater detector. Then, the craters identified are coupled to a database of known craters via a robust matching strategy leveraging geometric invariants [3] [18]. Finally, the camera pose is estimated by solving the related Perspective-n-Points (PnP) problem [20].

The deep learning-based crater detection module extends our previous work [4] to be deployed onboard. Comprehensive details about this module are presented in Section 4. By employing this selected detector, the system is capable of providing a comprehensive array of potential landmarks

for utilization in subsequent navigation stages. This DL-based detector has been trained using transfer learning on more than 800 real lunar monocular images obtained from the LRO cameras [4]. The crater detector has been validated on both WAC and NAC camera images, showing a remarkable mAP_{50} of above 0.7. In addition, the detector has been comprehensively tested for its robustness in the presence of noise, blur [4] and varying lighting conditions [38].

For each identified crater triad within the image, a set of camera-independent geometric features, namely interior angles and the ratio of crater radii to centroid distance, are calculated (See Figure 2). These descriptors effectively preserve the angular and spatial distribution of the crater triad. The same procedure is performed for the catalogue craters. A catalogue search window is identified based on the lander’s previous position, and assuming a nadir camera pointing. It is noteworthy that the number of the crater triads computed is approximately N^3 of the craters under investigation [3]. In the next phase, two lookup tables are constructed, one for the image craters and one for the catalogue craters, each including the unique identifier of each crater and the triad’s projective invariants. These tables are then merged using an inner join operation [3].

The output of the matching pipeline comprises a set of 2D-to-3D correspondences including the coordinates of the matched craters, both in the camera reference frame and in the Lunar Centered Lunar Fixed (LCLF) frame. To tackle the related PnP problem, and thus retrieve the camera pose in the LCLF frame, the Efficient Perspective-n-Point (EPnP) problem solver [19], has been employed. Its efficiency is particularly crucial in the resource-limited environment of a spacecraft, ensuring minimal computational burden. Ultimately, solving the PnP problem provides us with the monocular camera’s pose.

2.2. Relative Visual-based Navigation

As previously discussed in Section 2, the effectiveness of the crater-based navigation system diminishes as altitude decreases, given the reduction in visible catalogued craters. Hence, an alternative approach becomes necessary, which relies on tracking other salient surface features between consecutive frames [28] [37]. The proposed feature matching pipeline, integral to the implementation of the relative navigation algorithm, comprises two primary architectures: SuperPoint [6] and SuperGlue [32].

The SuperPoint network is a robust architecture for extracting feature points and descriptors from images. It utilizes a VGG-like [35] encoder to reduce the dimensions of the input image. Subsequently, two decoders jointly operate on this shared and spatially reduced representation of the input, outputting both the keypoints location and descriptors [6]. The network achieves lightness without sacrificing accuracy, making it suitable for tasks such as Simultane-

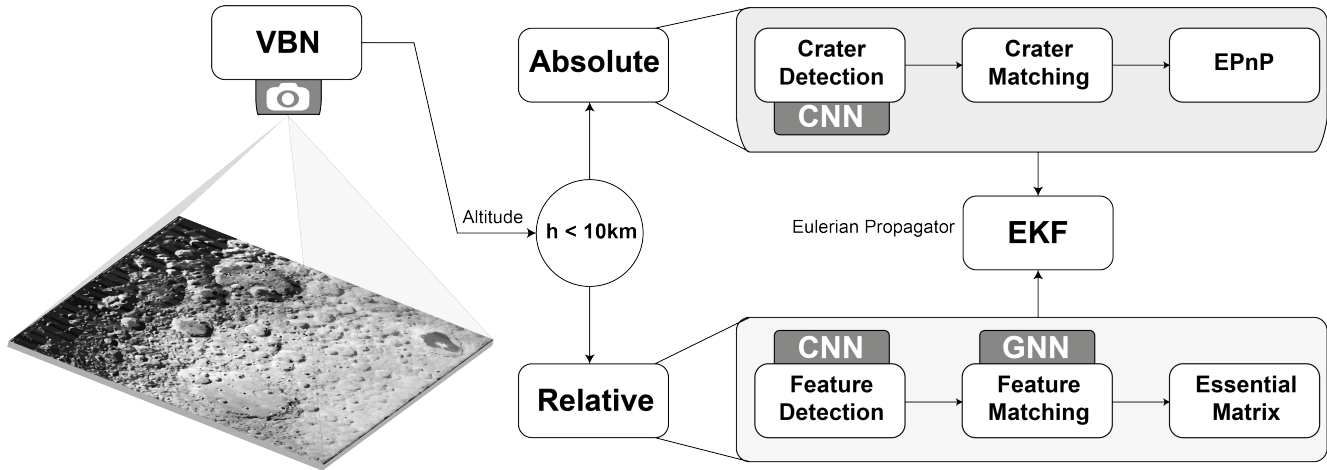


Figure 1. Flow diagram representing the proposed VBN system’s architecture in which the operational mode is selected based on specific criteria: if the altitude is below 10 km or no craters are detected, the relative mode is chosen. In absolute mode, images undergo crater detection and crater matches are identified. Then, the pose is estimated to solve the EpNP problem and integrated into the Extended Kalman Filter (EKF). For the relative mode, it processes two frames at a time and computes their offsets through an essential matrix obtained by feature tracking, then integrates these offsets into the EKF.




Figure 2. Example of selected projective invariants descriptors for a crater triad: interior angles and the ratio between crater radii and centroid distance.

ous Localization And Mapping (SLAM) and Structure from Motion (SfM) [5] [42] [21].

SuperGlue is a deep learning architecture that aims at matching sparse local features in images. Assignments are estimated by solving a differentiable optimal transport problem, whose costs are predicted by a Graph Neural Network (GNN). The architecture incorporates a flexible context aggregation mechanism based on attention, allowing it to re-

ason about the underlying 3D scene and feature assignments [32]. SuperGlue has been widely used in various applications, including SfM [21]. Recently, the LightGlue [22] architecture tackled the problem of sparse feature matching in a more computationally efficient fashion by dynamically reducing the number of layers depending on the difficulty of the input image pair, which varies based on the amount of visual overlap, appearance changes, or discriminative information. Furthermore, the network architecture discards at early stage points that are not matchable, thus focusing its attention on the visible area [22]. The relative pose between two consecutive viewpoints is retrieved from the collection of 2D point correspondences through essential matrix estimation, employing the RANSAC algorithm [13], followed by the matrix decomposition and check against the cheirality condition [27]. The scale ambiguity is addressed by incorporating the pipeline altimeter measurements generated within the virtual landing scenario.

3. Numerical Results

To evaluate and validate the performance of the proposed architecture, a testing pipeline has been established. Firstly, a dataset of synthetic images has been generated along a reference descent trajectory (See Section 3.1), to support the assessment of the VBN module’s performance across both operational modes. Subsequently, further testing of the crater detector on dedicated hardware has been performed to evaluate inference times and power consumption.

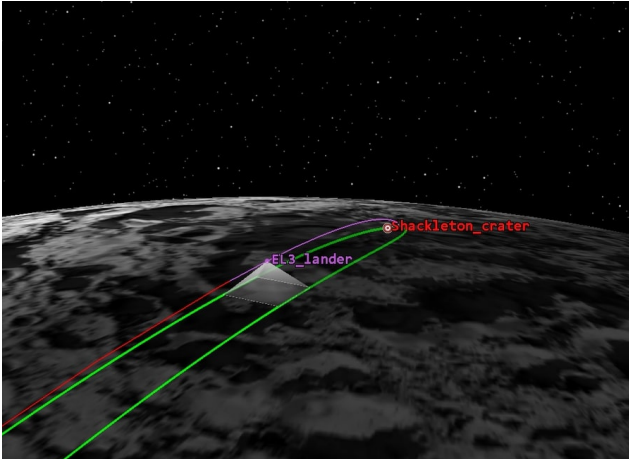


Figure 3. Reference descent trajectory considered for the EL3 lander [36], highlighted alongside the camera's FOV swath lines (in green). During the trajectory segment in red, the absolute module is employed, while the relative module is active in the purple segment.

3.1. Environment

The study examines a landing trajectory covering approximately 17 kilometers in altitude, with a duration of about 36 minutes. The chosen landing site is situated on the border of the Shackleton crater (Figure 3), positioned at the lunar South Pole. This selection was made based on analyses of the average annual exposure to solar illumination, aiming to identify the most viable locations for lunar landings [39]. While the investigation of optimal landing guidance falls outside the scope of this study, the reference trajectory was generated considering fuel consumption and thrust constraints, aligning with the specifications of the ESA Large Logistics Lander (EL3) [36].

The lander's dynamics were propagated factoring in the following forces acting on the spacecraft for both powered and ballistic segments: gravitational force from the main celestial body, third body effects from both the Sun and the Earth, and solar radiation pressure on the spacecraft surface at a distance of 1 AU from the Sun. The lander has an initial wet mass of 2500 kg and a dry mass of 1200 kg.

A dataset of synthetic images of the lunar surface was generated starting from the DEM generated by the LRO LOLA. With a resolution up to 5 m/pixel, this DEM enabled to construct a high fidelity 3D model of the lunar terrain traversed by the lander. The surface was subsequently imported into Unreal Engine 5 [11], leveraging its capabilities to handle high-polygon meshes and high-fidelity scenes through innovative technologies such as Nanite [10], Virtual Shadow Maps (VSMs) [12] and Lumen [9]. This virtual landing scenario facilitated the creation of high-resolution images under varying lighting conditions. Figure 4 depicts the lan-

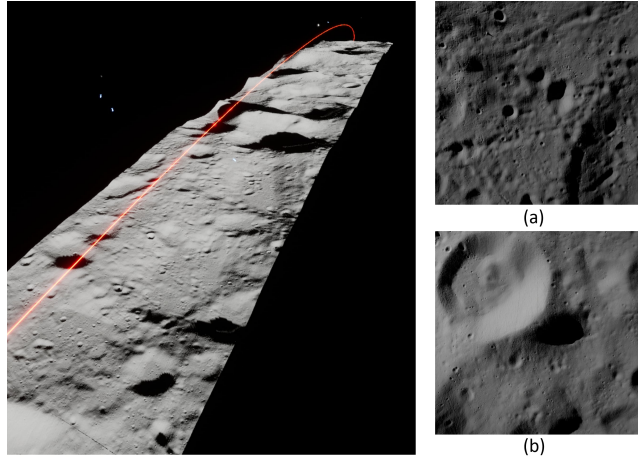


Figure 4. Reference descent trajectory within the virtual landing scenario based on Unreal Engine [11] (left). The synthetic image (a) has been captured from the environment at longitude 0.96° W and latitude 75.80° S. Image (b) has been generated at longitude 1.66° W and latitude 82.62° S

der's trajectory within the virtual scenario, alongside two illustrative examples of synthetic images generated from this environment. A material simulating lunar regolith was applied to the surface to enhance realism. Post-processing techniques, such as Contrast Limited Adaptive Histogram Equalization (CLAHE), were employed to ensure alignment between the synthetic images dataset and a dataset of real images captured by the LRO's Wide Angle Camera (WAC).

3.2. Experiments

Both the absolute and relative navigation algorithms underwent testing using a sequence of synthetic images generated from the virtual landing scenario in Unreal Engine. These images were captured along the reference trajectory described in Section 3.1. Firstly, the accuracy of the position estimation was assessed for the absolute operational mode. Concerning the relative navigation mode, two investigations were conducted: one focusing on estimating the displacement error relative to terrain, and the other examining the relative rotation error along a reference maneuver. A nadir-pointing camera with a resolution of 1024×1024 and with a FOV of 90° has been assumed throughout the analyses.

3.2.1 Absolute Mode

The performance assessment of the absolute pose estimation algorithm was conducted according to the following methodology: a sequence of uniformly spaced images was captured along the first portion of the landing trajectory, at 15 km of altitude. Each image was input to the absolute pose estimation algorithm, which performed crater detec-

tion, matching with the selected catalogue [30] and pose retrieval by solving the related PnP problem through the EPnP solver. The camera frame rate has been set to 1/5 Hz.

Figure 5 provides the outcomes of the crater matching task for a synthetic image sample that the crater detector has not encountered during the training phase. The detected craters (in green) have been matched with the corresponding catalogue craters (in yellow).

The average absolute position error has been computed as the Euclidean norm of the vector difference between estimated and true lander positions in the LCLF reference frame. This mean error has been estimated to be 270.1 m.

For all stages of the absolute navigation algorithm, computational times were computed using commercial hardware (Intel® Core™ i7-8086K Processor @5GHz and 16 GB of system RAM) yielding the following results: approximately 1.08 s on average for the crater detector CPU inference, about 90 ms for the matching phase, and less than 2 ms for the EPnP solver.




Figure 5. An example of crater detection and matching. The green craters are those detected by the crater detector. The yellow craters are the corresponding catalogue craters.

3.2.2 Relative Mode

The relative navigation algorithm showcased in Section 2.2 has been tested along the final segment of the trajectory, covering the altitude range from 10 km to 2 km and lasting approximately 150 seconds. To assess the pipeline’s navigation performance, the displacement error relative to the ground has been evaluated. Figure 6(a) display said error, split into its horizontal and vertical components, and eval-

uated at each new frame acquisition along the selected trajectory portion. The camera’s framerate was set to 1/4 Hz, and the number of keypoints to detect for each frame was configured to 1024 for the testing procedure. Both the horizontal and vertical errors display a negative trend with decreasing altitude. The average displacement error along the trajectory has been estimated to be approximately 27.4 m for the horizontal component. the vertical error component is about 0.8 meters and can be attributed to the altimeter measurements being directly integrated into the navigation pipeline.

The rotation error relative to terrain was estimated by simulating a yaw slew maneuver along the final trajectory segment, and by evaluating the error committed on the measured rotation. The chosen maneuver starts with a 30° rotation lasting for about 20 seconds, followed by a 40 seconds coasting phase where the lander proceeds along the trajectory in a crabbed fashion. Finally, the maneuver ends with the lander performing another 30° slew maneuver to recover the initial attitude configuration.

By exploiting the virtual scenario capabilities, it was possible to generate a synthetic image sequence during the lander’s slew at varying camera frame rates. This sequence was subsequently used in the pose estimation algorithm to retrieve the instantaneous and mean relative attitude errors. Figure 6(b) depicts a comparison between the actual rotation and the rotation measured by the relative navigation pipeline. Very limited deviation from ground truth can be appreciated, with a maximum cumulated error of about 0.7°, indicating a good correlation between the rotation estimates provided by the pipeline and the actual rotations.

Testing the relative navigation algorithm on commercial hardware (see Section 3.2.1) yielded an average LightGlue CPU inference time of 3.81 s, and 50 ms for the pose estimation.

To compare the relative efficiency and effectiveness of each operational mode in VBN, Table 1 presents an analysis contrasting the positional accuracy of both the absolute and relative modes. The analysis highlights that the absolute navigation pipeline yields an estimation error of approximately 300 m along the trajectory, consistent with the anticipated performance at these relevant altitudes. The relative mode demonstrates notably low error margins in horizontal displacement, consistently remaining under 150 m.

4. Toward Onboard-AI Implementation

As stated above, with the expanding uptake of AI onboard satellites, the possibility of effectively streamlining computational loads to AI-dedicated chips is increasing. These components can accelerate or execute with high energy efficiency in some modules of the processing chain. The landscape of platforms for onboard AI deployment in satellites is notably limited to a few devices supported by various




Figure 6. (a) Horizontal and vertical components of the relative displacement error along the selected trajectory segment. (b) Comparison between actual rotation and rotation measured by the relative navigation pipeline. The maximum deviation is approximately 0.7°.

VBN Mode	Performance
Absolute	Mean Position Error: 270.1 m @ 1/5Hz
	Mean Horizontal Position Error: 216.1 m @ 1/5Hz
	Mean Vertical Position Error: 54.0 m @ 1/5Hz
Relative	Mean Horizontal Displacement Error: 27.4 m @ 1/4Hz
	Mean Vertical Displacement Error: 0.8 m @ 1/4Hz

Table 1. Absolute and relative mode accuracies in terms of mean position error, mean horizontal displacement and mean vertical displacement.

In-Orbit Demonstration on Class V/IV² spacecrafts and on-ground radiation testing. This can be explained by the complexity of qualifying ground platforms for space applications and by the low maturity of AI models and workflows for onboard deployment.

In this study, we show a demonstration of the optimization and testing of the crater detector module. To facilitate a rigorous and exhaustive evaluation, we conducted empirical testing of the optimized crater detector module on two established hardware platforms with a proven flight legacy.

²ESA Mission Classification and Project Adoption of New Microelectronics Development, available online at: <https://tinyurl.com/2s5shdkj>

Specifically, our selection included the Intel® Neural Compute Stick 2 (Intel® NCS2) and the Ubotica CogniSAT-XE1, which incorporates the Intel® Movidius™ Myriad™ 2 chip.

A simple but effective model for the crater detection module has been tested: we combine ResNet-50 as a backbone for feature extraction with head layers similar to the ones of SSD [23]. To tailor the model more effectively to on-board devices, we removed certain layers³ from the original implementation for a more streamlined design. A pivotal modification in the *conv4_x* layer involved setting strides to a uniform 1 × 1, diverging from typical configurations. *BatchNorm* layers were added after each convolution in the SSD heads to boost stability and training efficiency.

The model training has been performed with mixed precision using an NVIDIA A100 GPU. Stochastic Gradient Descent (SGD) with a momentum of 0.9 was utilized for optimization. A critical element in our training methodology was the learning rate scheduler. We adopted a cosine annealing approach, starting with an initial learning rate of 2.6e-3. To enhance training, we implemented learning rate decay, and a linear warmup strategy during the first epoch, incrementally increasing the learning rate from a low to a higher value. This technique helps avert early training divergence due to an excessively high starting learning rate. Subsequently, the model underwent post-training quantization using the OpenVINO toolkit (v2020.3) to compile it in half-precision. These specific training strategies were crucial in optimizing the model for quantization, ensuring

³All the *conv5_x*, the *Average Pooling*, and *fully-connected* layers in the classification heads.

Metric	IoU	Area	maxDets	Value
<i>AP</i>	0.50:0.95	<i>all</i>	100	0.34
			0.50	0.67
			0.75	0.30
	0.50:0.95	<i>small</i>		0.23
			<i>medium</i>	0.55
			<i>large</i>	0.58
<i>AR</i>	0.50:0.95	<i>all</i>	1	0.047
			10	0.30
			100	0.46
		<i>small</i>		0.38
			<i>medium</i>	0.65
			<i>large</i>	0.66

Table 2. Average Precision (AP) and Average Recall (AR) prompted at different scales and intervals of Intersection over Union (IoU).

both effective performance and compatibility. The model’s efficacy was assessed using metrics like Average Precision (AP) and Average Recall (AR), considering various scales and Intersection over Union (IoU) thresholds. Table 2 presents the ultimate results obtained through an exhaustive hyper-parameter tuning process, employing a grid search technique. Ultimately, the model was successfully deployed on two designated devices, yielding inference times under one second. Figure 7 presents a bar plot comparing the inference times for processing a single image (tensor size $1 \times 3 \times 300 \times 300$) on each device. As it can be observed, the NCS2 demonstrates a performance that is three times faster than that of the XE1 board.

Nonetheless, despite the achievements of our study, it is important to acknowledge certain limitations and identify areas for further improvement. One of the primary areas of focus in our ongoing research is the optimization of the




Figure 7. Latency comparison of the quantized model for crater detection. The bar plot illustrates the time performance, in seconds, of two different devices under test, i.e., the CogniSAT-XE1 (XE1) and the Intel® Neural Compute Stick 2 (NCS2).

LighGlue network. Currently, the network exhibits promising capabilities; however, its optimization for onboard deployment remains a critical challenge. This requires a careful balance between maintaining high accuracy and minimizing computational demands. Additionally, there is a significant opportunity to enhance the efficiency of the crater matching pipeline. As it stands, the pipeline faces a computational bottleneck, primarily due to its quadratic increase in computational complexity relative to the number of craters [3]. This issue poses a considerable challenge, especially in scenarios involving large datasets with numerous craters. To address this, we aim to develop more sophisticated algorithms that can reduce the computational load without compromising the accuracy of crater matching.

One potential approach is to implement more efficient data structures or to leverage advanced machine learning techniques that can intelligently prioritize or filter crater comparisons. Another avenue could be the integration of parallel processing techniques or hardware accelerators like VPUs, which can significantly speed up computations. Furthermore, the current quadratic complexity makes scalability a concern, particularly for extensive lunar or planetary surveys where thousands of craters need to be analyzed. Streamlining the crater-matching process would not only improve the speed but also enhance the overall feasibility of large-scale applications.

In addition, the optimization of the LighGlue network for onboard use and the improvement of the crater matching pipeline to tackle computational inefficiencies stand out as critical areas for future research. These enhancements will not only augment the practicality of our study but also broaden its applicability in various space exploration and remote sensing missions.

5. Conclusion

This research introduces a tailored solution for lunar landing operations without GNSS signal assistance. The study innovates by integrating a dual-mode VBN system, adept at switching between absolute and relative modes depending on altitude and crater visibility. Our simulated descent experiments over Shackleton crater mark a first step in validating the AI-based system for onboard use. In terms of accuracy, the system’s relative mode achieved an average horizontal displacement error of merely 27.4 m at 1/4 Hz and an impressively low vertical displacement error of 0.8 m at the same frequency, showcasing the system’s high precision. Our evaluations using CogniSAT-XE1 and Neural Compute Stick 2 have confirmed the capability of operating the crater detection module with inference times under one second and with an overall AP of 0.67. Future research will focus on tuning the feature matching network for onboard deployment and enabling the crater matching process, critical aspects for enabling the practicality and expanding the potential applications of this study in lunar exploration missions.

References

- [1] Sivert Bakken, Marie B Henriksen, Roger Birkeland, Dennis D Langer, Adriënne E Oudijk, Simen Berg, Yeshe Pursley, Joseph L Garrett, Fredrik Gran-Jansen, Evelyn Honoré-Livermore, et al. Hypso-1 cubesat: first images and in-orbit characterization. *Remote Sensing*, 15(3):755, 2023. 2
- [2] John A Christian, Harm Derksen, and Ryan Watkins. Lunar crater identification in digital images. *The Journal of the Astronautical Sciences*, 68(4):1056–1144, 2021. 1
- [3] Roberto Del Prete and Alfredo Renga. A Novel Visual-Based Terrain Relative Navigation System for Planetary Applications Based on Mask R-CNN and Projective Invariants. *Aerotecnica Missili & Spazio*, 101(4):335–349, 2022. 1, 3, 8
- [4] Roberto Del Prete, Alfonso Saveriano, and Alfredo Renga. A deep learning-based crater detector for autonomous vision-based spacecraft navigation. In *2022 IEEE 9th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, pages 231–236, 2022. 3
- [5] Chengqi Deng, Kaitao Qiu, Rong Xiong, and Chunlin Zhou. Comparative study of deep learning based features in slam. In *2019 4th Asia-Pacific Conference on Intelligent Robot Systems (ACIRS)*, pages 250–254. IEEE, 2019. 4
- [6] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabbinovich. Superpoint: Self-supervised interest point detection and description. In *CVPR Deep Learning for Visual SLAM Workshop*, 2018. 3
- [7] Lena M Downes, Ted J Steiner, and Jonathan P How. Lunar terrain relative navigation using a convolutional neural network for visual crater detection. In *2020 American Control Conference (ACC)*, pages 4448–4453. IEEE, 2020. 1
- [8] Lena M Downes, Ted J Steiner, and Jonathan P How. Lunar terrain relative navigation using a convolutional neural network for visual crater detection. In *2020 American Control Conference (ACC)*, pages 4448–4453. IEEE, 2020. 2, 3
- [9] Inc. Epic Games. Lumen global illumination and reflections, . <https://docs.unrealengine.com/5.0/en-US/lumen-global-illumination-and-reflections-in-unreal-engine/> [Accessed: 03/03/2024]. 5
- [10] Inc. Epic Games. Nanite virtualized geometry, . <https://docs.unrealengine.com/5.0/en-US/nanite-virtualized-geometry-in-unreal-engine/> [Accessed: 03/03/2024]. 5
- [11] Inc. Epic Games. Unreal engine 5, . <https://www.unrealengine.com/en-US/unreal-engine-5> [Accessed: 03/03/2024]. 5
- [12] Inc. Epic Games. Virtual shadow maps, . <https://docs.unrealengine.com/5.0/en-US/virtual-shadow-maps-in-unreal-engine/> [Accessed: 03/03/2024]. 5
- [13] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24:381–395, 1981. 4
- [14] Gonzalo Mateo García. *Transfer learning of deep learning models for cloud masking in optical satellite images*. PhD thesis, Universitat de València, 2023. 2
- [15] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Ver-cruyssen, et al. The ϕ -sat-1 mission: The first on-board deep neural network demonstrator for satellite earth observation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2021. 2
- [16] Andrew E Johnson and James F Montgomery. Overview of terrain relative navigation approaches for precise lunar landing. In *2008 IEEE Aerospace Conference*, pages 1–10. IEEE, 2008. 1
- [17] Shreeyam Kacker, Alex Meredith, Kerri Cahoy, and Georges Labreche. Machine learning image processing algorithms onboard ops-sat. 2022. 2
- [18] Kiduck Kim, Youeyun Jung, and Hyochoong Bang. *Projective invariant based crater matching for visual navigation in planetary landing*. 3
- [19] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Eppn: An accurate o(n) solution to the pnp problem. *International Journal of Computer Vision*, 81(2): 155–166, 2008. 3
- [20] Shiqi Li, Chi Xu, and Ming Xie. A robust o(n) solution to the perspective-n-point problem. *IEEE transactions on pattern analysis and machine intelligence*, 34, 2012. 3
- [21] Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson, and Marc Pollefeys. Pixel-Perfect Structure-from-Motion with Featuremetric Refinement. In *ICCV*, 2021. 4
- [22] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. 4
- [23] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016. 2, 7

- [24] Iain Martin, Martin Dunstan, and Manuel Sanchez Gestido. Planetary surface image generation for testing future space missions with pangu. In *2nd RPI Space Imaging Workshop*. Sensing, Estimation, and Automation Laboratory, 2019. 2
- [25] Gonzalo Mateo-García, Josh Veitch-Michaelis, Cormac Purcell, Nicolas Longepe, Simon Reid, Alice Anlind, Fredrik Bruhn, James Parr, and Pierre Philippe Mathieu. In-orbit demonstration of a re-trainable machine learning payload for processing optical imagery. *Scientific Reports*, 13(1):10391, 2023. 2
- [26] Shervin Minaee, Yuri Boykov, Fatih Porikli, Antonio Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7):3523–3542, 2022. 2
- [27] D. Nister. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 2004. 4
- [28] Luca Ostrogovich. Study of ai guided visual based navigation models for autonomous landing on lunar surface scenarios. M.sc, University of Naples Federico II, Italy, 2023. 3
- [29] Leo Pauly, Wassim Rharbaoui, Carl Shneider, Arunkumar Rathinam, Vincent Gaudillière, and Djamila Aouada. A survey on deep learning-based monocular spacecraft pose estimation: Current state, limitations and prospects. *Acta Astronautica*, 212:339–360, 2023. 2
- [30] Stuart J. Robbins. A New Global Database of Lunar Impact Craters > 1-2 km: 1. Crater Locations and Sizes, Comparisons With Published Databases, and Global Analysis. *Journal of Geophysical Research (Planets)*, 124(4):871–892, 2019. 3, 6
- [31] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 2
- [32] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 3, 4
- [33] Ari Silburt, Chenchong Zhu, Mohamad Ali-Dib, Kristen Menou, and Alan Jackson. Deepmoon: Convolutional neural network trainer to identify moon craters. *Astrophysics Source Code Library*, pages ascl-1805, 2018. 2, 3
- [34] Stefano Silvestrini, Margherita Piccinin, Giovanni Zanotti, Andrea Brandonisio, Ilaria Bloise, Lorenzo Feruglio, Paolo Lunghi, Michele Lavagna, and Mattia Varile. Optical navigation for lunar landing based on convolutional neural network crater detector. *Aerospace Science and Technology*, 123:107503, 2022. 2, 3
- [35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. 3
- [36] the European Space Agency. Esa - argonaut. https://www.esa.int/Science_Exploration/Human_and_Robotic_Exploration/Exploration/Argonaut [Accessed: 03/05/2024]. 5
- [37] Giuseppe Tomasicchio, Luca Andolfi, Marco Brancati, Arsenio Maria Di Donna, Simone Giannattasio, Roberto Del Prete, Luca Ostrogovich, Alfredo Renga, Michele Grassi, Michele Ceresoli, et al. Multi-sensor fusion strategy for improved navigation in lunar landing missions. Presented at the ASI workshop on GNC technologies for accessing the lunar surface, Rome, Italy, November 2023. 3
- [38] Giuseppe Tomasicchio, Luca Andolfi, Marco Brancati, Arsenio Maria Di Donna, Simone Giannattasio, Roberto Del Prete, Luca Ostrogovich, Alfredo Renga, Michele Grassi, Michele Ceresoli, et al. Multi-sensor fusion for improved navigation in lunar landing missions. In *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, pages 3549–3576, 2023. 3
- [39] Kathryn Hambleton Vanessa Lloyd, Alana Johnson. Nasa identifies candidate regions for landing next americans on moon, 2022. <https://www.nasa.gov/news-release/nasa-identifies-candidate-regions-for-landing-next-americans-on-moon/> [Accessed: 03/05/2024]. 5
- [40] Hao Wang, Jie Jiang, and Guangjun Zhang. Crateridnet: An end-to-end fully convolutional neural network for crater detection and identification in remotely sensed planetary images. *Remote Sensing*, 10:1067, 2018. 2, 3
- [41] Wei Wang, Yujing Yang, Xin Wang, Weizheng Wang, and Ji Li. Development of convolutional neural network and its application in image classification: a survey. *Optical Engineering*, 58:040901 – 040901, 2019. 2
- [42] Zhenyu Yin, Dan Feng, Chao Fan, Chengen Ju, and Feiqing Zhang. Sp-vslam: monocular visual-slam algorithm based on superpoint network. In *2023 15th International Conference on Communication Software and Networks (ICCSN)*, pages 456–459, 2023. 4