A Hybrid AI-assisted Hazard Detection and Avoidance Pipeline for Autonomous Lunar Landing

SAPIENZA UNIVERSITÀ DI ROMA

Mohamed El Awag[1*], Riccardo Teodori[1], Ludovica Cavalieri[1], Antonio Genova[1]

[1] Sapienza University of Rome. Via Eudossiana 18, Roma, 00184, Italy

[*] mohamed.elawag@uniroma1.it

1. Overview

Safe **autonomous landing** is a key enabling capability for next-generation lunar missions [1]. Conventional Guidance, Navigation, and Control architectures show limited robustness and **real-time responsiveness** in such unstructured environments [2].
We propose a **hybrid AI-assisted** pipeline that integrates data-driven models with deterministic control to address both Hazard Detection (HD) and Hazard Avoidance (HA). Computationally intensive operations are shifted to learning-based models, where **complexity is handled offline** during training, enabling real-time execution during terminal descent. The framework combines Convolutional Neural Networks for **vision-based terrain analysis**, a Multi-Layer Perceptron (MLP) for **rapid trajectory replanning**, and a Model Predictive Control (MPC) layer for **robust trajectory tracking**.

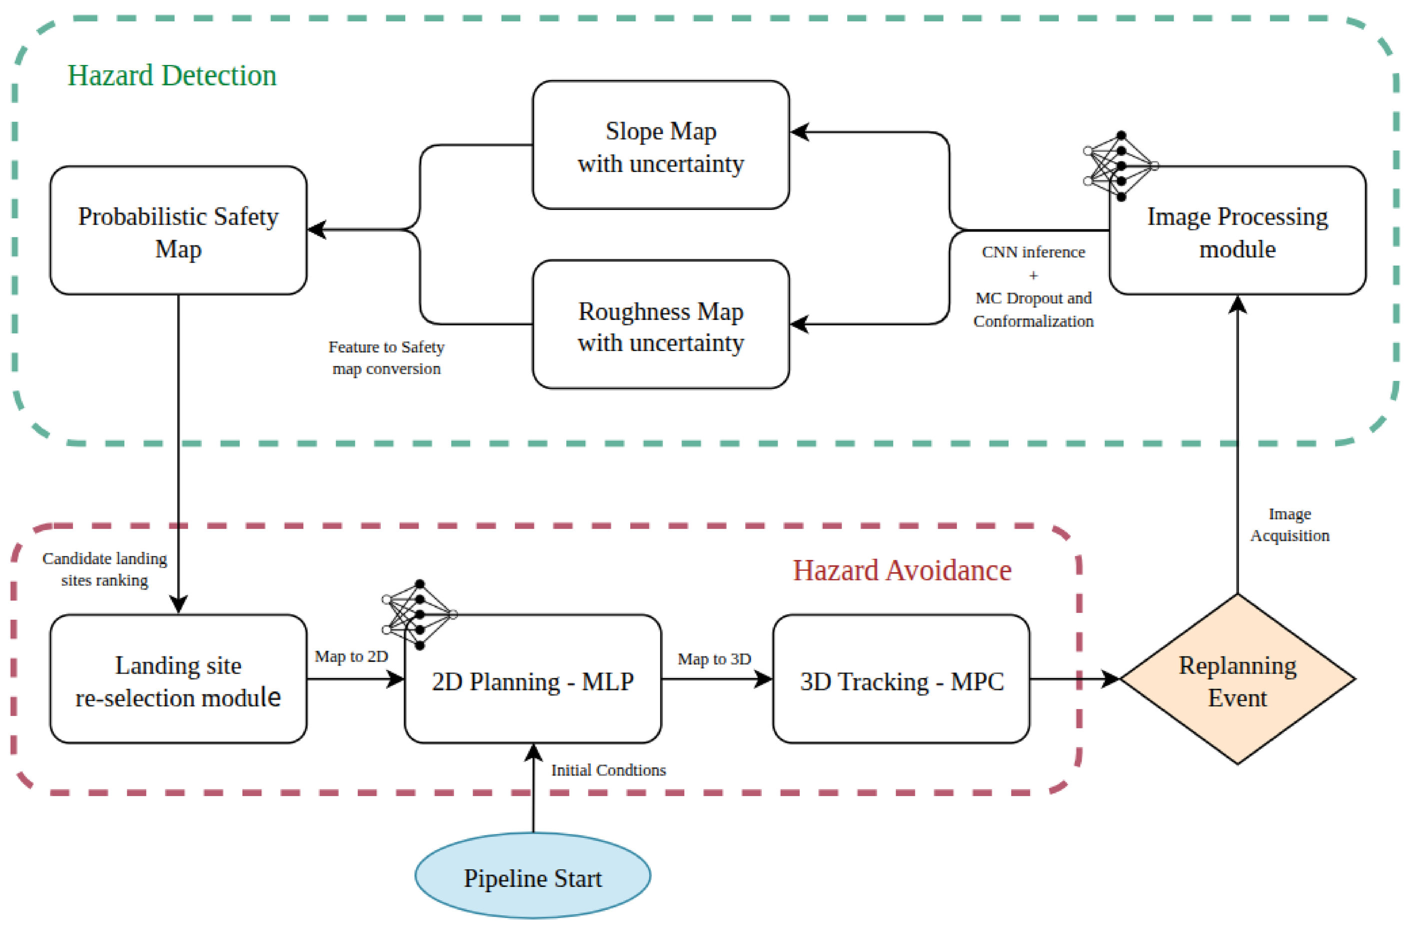

Fig. 1. *Schematic representation of the proposed HDA Pipeline. The modules integrating AI elements are highlighted with a network icon.*

2. Hazard Detection and Landing Site Re-selection

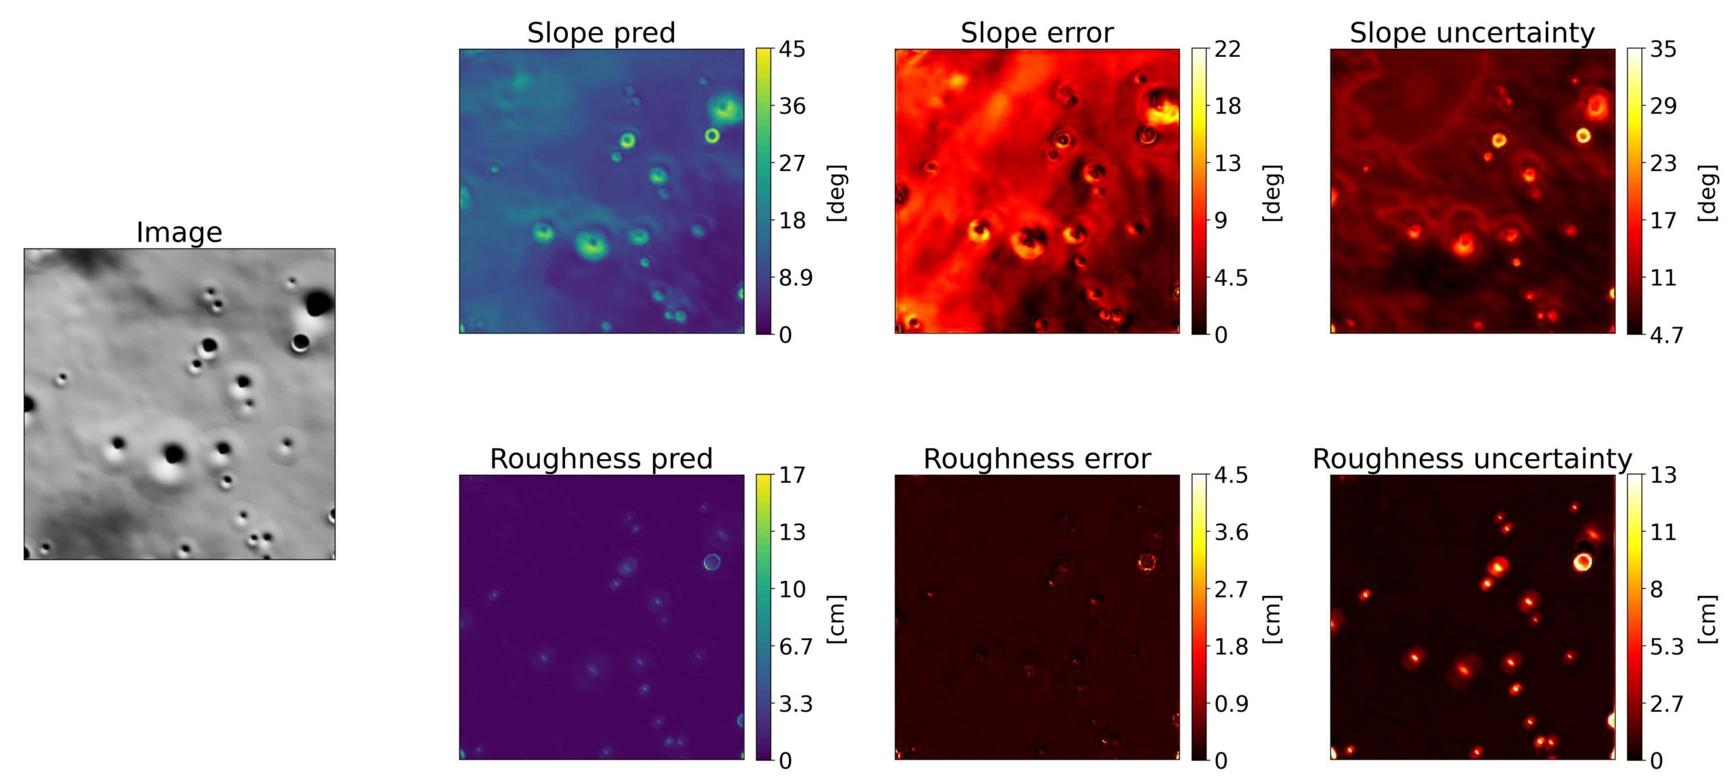

Fig. 2. *Preliminary results of networks prediction for slope and roughness with associated errors and uncertainties. The input image is taken from 500 m of altitude by a nadir-pointing camera. Although conservative, the uncertainty provides a consistent representation of networks error.*

Predictions and uncertainties are fused into a **probabilistic safety map** [9], enabling risk-aware evaluation of candidate landing sites. Candidates are selected through a **search-and-ranking procedure** based on the Vehicle Dispersion Footprint Ellipse [10], which accounts for the lander's physical dimensions. Sites below a 50% safety threshold are discarded, and the remaining ones are ranked by a **composite score** combining **safety** probability, **maneuver cost** (Euclidean distance from the current position), and a **clearance** score measuring distance from hazardous regions.

Monocular grayscale descent images are processed by **two separate U-Nets** [3] that perform **dense regression** of continuous **slope** and surface **roughness** maps, overcoming the limited resolution of classification-based methods [4]. The networks are trained on a **synthetic dataset** generated through a Blender-based [5] photorealistic pipeline, with ground-truth labels computed via spatial filtering [6]. Pixel-wise **uncertainty** is estimated through **Monte Carlo Dropout** [7] and then calibrated via **Conformal Prediction** [8], yielding conservative, distribution-free bounds for safety-critical decisions.

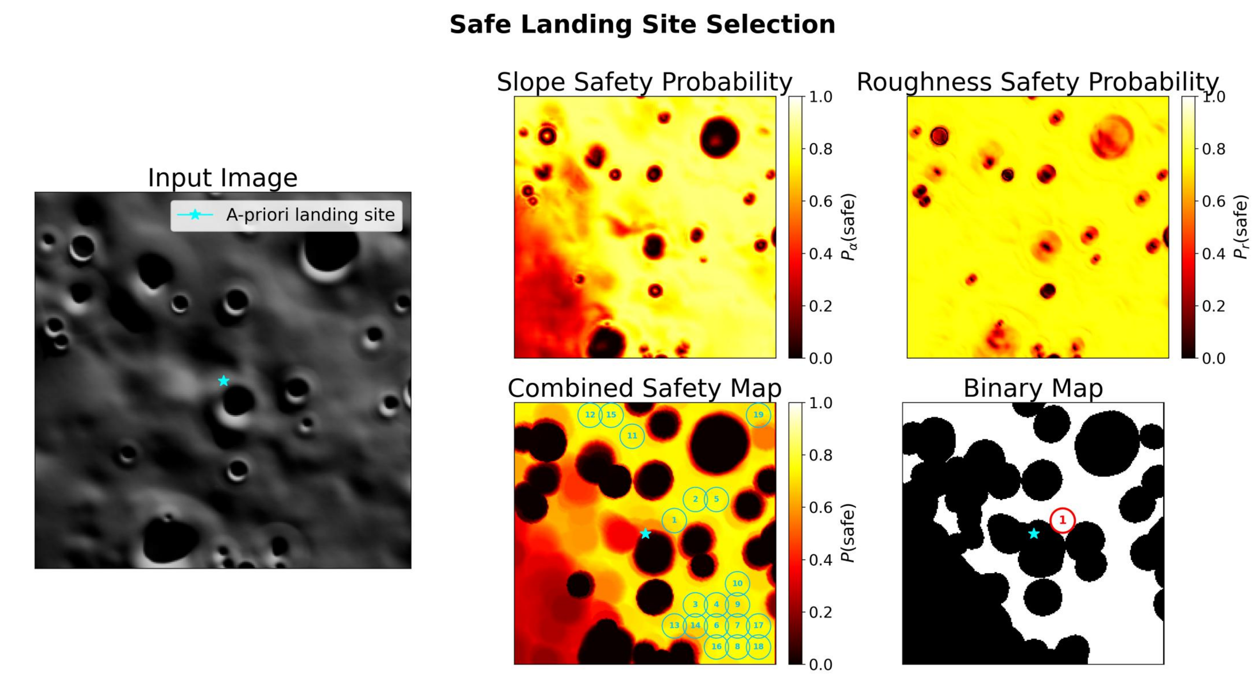

Fig. 3. *Safety map generation and landing site re-selection for a landing image acquired from 500 m of altitude.*

3. Hazard Avoidance

Once a safe site is selected, the HA module **computes a trajectory** from the current state to the target. To avoid the cost of onboard numerical control optimization, a **MLP is trained offline** to approximate the solution of **a mass-optimal control problem** derived from Pontryagin's Minimum Principle. The dataset comprises roughly 10,000 fuel-optimal soft-landing trajectories, generated via **indirect optimization** with homotopy techniques and **sampled into state-action pairs** [11]. Once trained, the MLP returns solutions with tens-of-milliseconds latency. A polynomial-guidance fallback handles **out-of-distribution initial conditions**, prioritizing constraint satisfaction and mission safety [12, 13]. **Planning** is performed in a reduced **2D** manifold and the resulting trajectory is **mapped to 3D** as a reference for tracking.

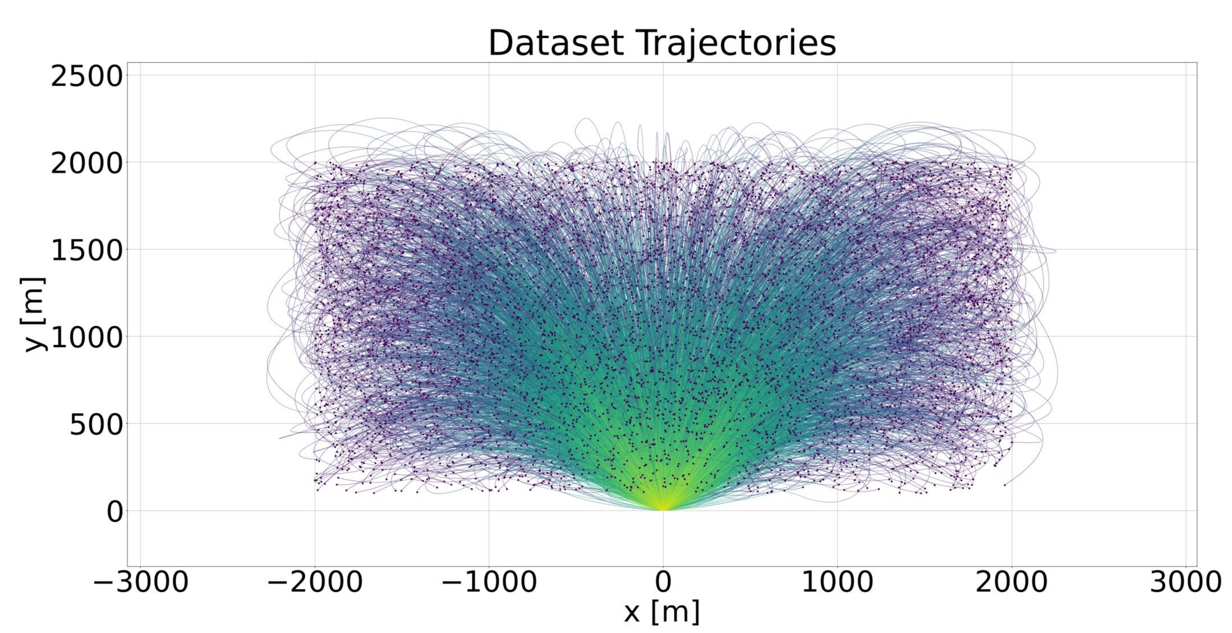

Fig. 4. *Visualization of the MLP training dataset showcasing the different optimal landing trajectories.*

4. Tracking

A **Model Predictive Control** layer tracks the MLP-derived reference within a full 3D dynamical framework, solving a finite-horizon optimal control problem in receding-horizon fashion. The cost penalizes **tracking error** along the horizon, while a terminal term enforces **soft landing** conditions on position, velocity, and attitude. The MPC **compensates for disturbances** (*e.g.*, residual out-of-plane accelerations after the replanning event) while enforcing actuator and system constraints. A preliminary Monte Carlo analysis confirms the **robustness** of the tracking layer, achieving sub-meter level touchdown accuracy and near-zero touchdown velocity both in vertical and horizontal components. Tracking performance were evaluated considering perfect localization.

5. Summary

The proposed pipeline was implemented in a **ROS2** environment to validate the **real-time** feasibility of its modules and their integrated, closed-loop operation during terminal descent. The hybrid AI-assisted HDA framework combines deep learning, uncertainty quantification, and robust control to **overcome key limitations** of existing approaches, enabling high-precision, real-time landing in challenging lunar environments. Future work targets validation in **hardware-in-the-loop** environments and deployment in real mission scenarios.

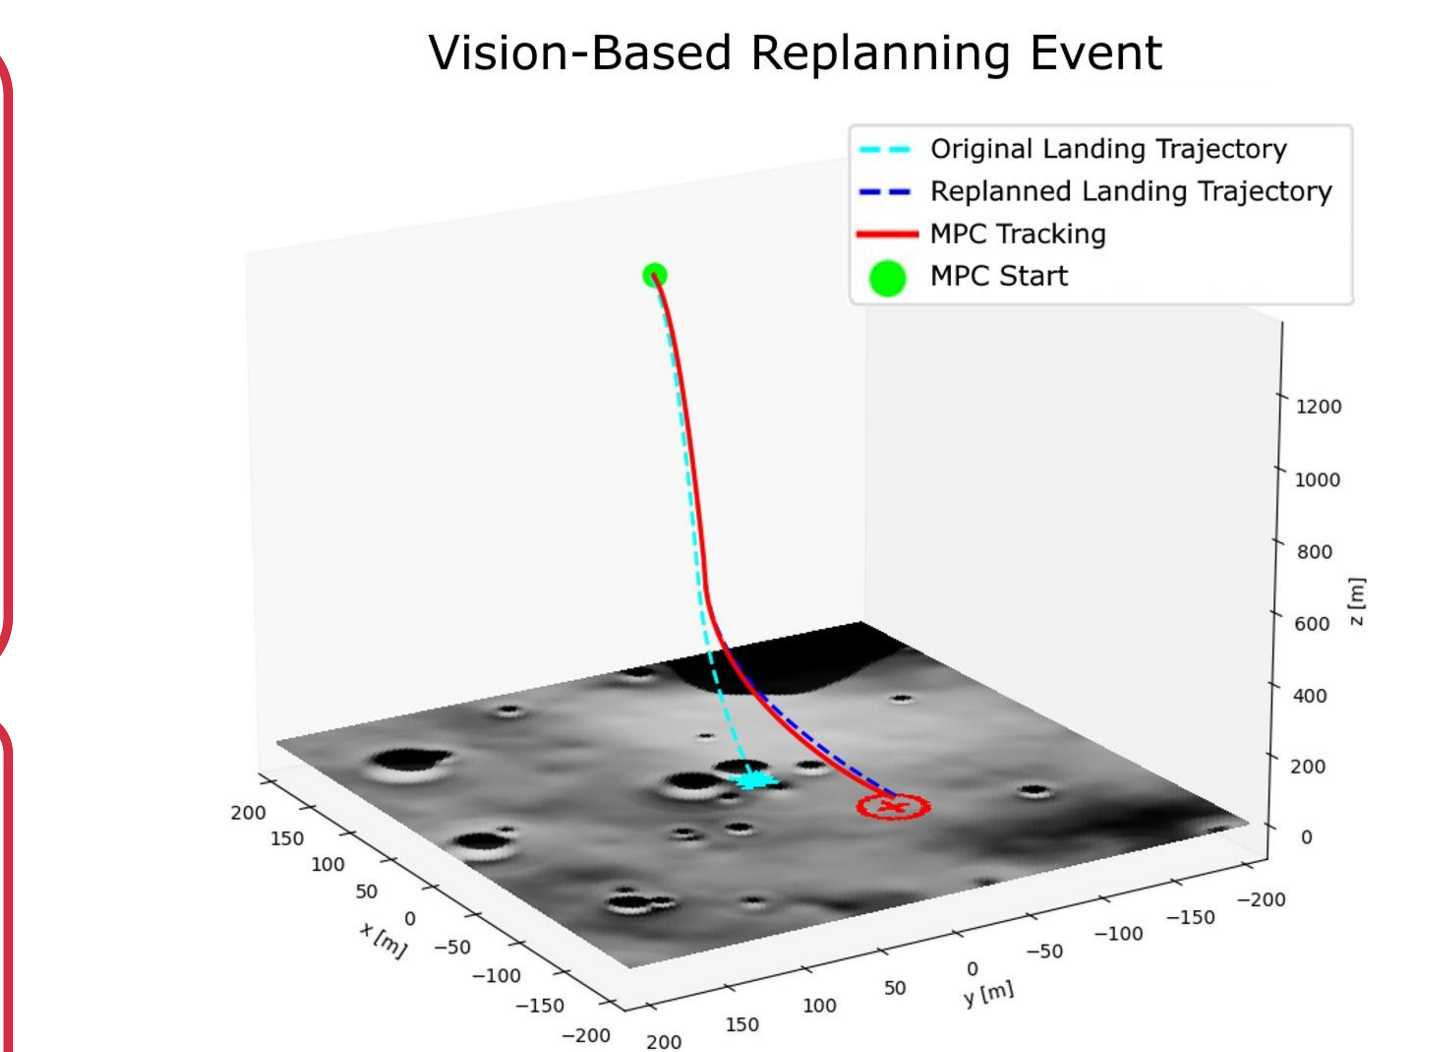

Fig. 5. *Integration of the HDA pipeline: the HD module selects a new landing site, the HA module replans the trajectory (cyan to blue), and the MPC tracks the updated reference (red), guiding the lander to the final target.*

References:
[1] J. Schonfeld, IEEE SMC, 2023. [2] C. D. Epp et al., IEEE Aerospace Conf., 2007. [3] O. Ronneberger et al., MICCAI, 2015. [4] M. El Awag et al., IEEE Aerospace Conf., 2026. [5] Blender Online Community, 2018. [6] L. Cavalieri et al., AIAA SciTech, 2026. [7] K. Tomita et al., J. Spacecraft and Rockets, 2022. [8] A. N. Angelopoulos and S. Bates, arXiv:2107.00363, 2022. [9] T. Ivanov et al., AIAA GNC Conf., 2013. [10] Y. Jin et al., Transactions in GIS, 2025. [11] C. Sanchez-Sanchez and D. Izzo, J. Guidance, Control, and Dynamics, 2018. [12] E. Wong et al., AIAA Atmospheric Flight Mechanics Conf., 2002. [13] S. Ploen et al., AIAA/AAS Astrodynamics Specialist Conf., 2006.

ICRA IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION 2026 VIENNA

IEEE International Conference on Robotics & Automation Vienna, Austria, 1–5 June 2026