POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Receding-Horizon Guidance and Control Around an Unknown Non-Cooperative Spacecraft for Deep Learning Based 3D Reconstruction

Laurea Magistrale in Space Engineering - Ingegneria Industriale

Author: **DAVIDE BELLINI**

Advisor: **STEFANO SILVESTRINI**

Co-advisor: **MADHUR TIWARI**

Academic year: **2024-2025**

1. Introduction

Autonomous close-proximity operations (CPO) are gaining relevance across several space mission domains, including in-orbit servicing, end-of-life management, active debris removal and the inspection of uncooperative targets. These scenarios require safe operation in the immediate vicinity of another object without cooperative aids, under tight margins and limited actuation authority, while remaining robust to navigation errors and model mismatch. Similar needs arise in small-body exploration, where sustained short range observation supports the estimation of shape, rotation state and surface properties.

Imaging is central to CPO, enabling optical navigation and non-contact target characterization from appearance, relative motion, and geometric signatures. Image sequences can also support the estimation of target pose, angular motion, and geometry. Recent trends leverage multi-view imagery to build three-dimensional target models. Learning based reconstruction can fuse images into continuous scene representations, yet it remains sensitive to data quality. This motivates online data acquisition strategies that shape the guidance trajectory to produce compatible datasets.

2. Problem Statement and Main Contributions

The goal is to enable autonomous inspection of an unknown, non-cooperative target by generating an image dataset suitable for deep learning based 3D reconstruction. The target orbital motion is assumed to be known, while geometry and surface properties are not available a priori. As a result, inspection cannot be precomputed. It must be executed online, maintaining safe relative motion and admissible imaging conditions while adapting future viewpoints to the evolving observation history.

This thesis proposes an inspection-driven onboard architecture with the following contributions:

- a mapping strategy to quantify inspection progress and promote uniform viewpoint distribution;
- an online replanning logic that updates exploration directions as coverage evolves;
- closed-loop receding-horizon guidance demonstrating safe proximity motion and sustained imaging capability.

3. Methodology

To enable autonomous inspection for reconstruction oriented mapping, the guidance strategy is structured as a closed loop driven by imaging campaign progress. The system alternates between selecting an exploration direction that prioritizes under observed viewing regions and generating a feasible 6-DoF trajectory that safely tracks it while maintaining admissible imaging conditions. Progress is monitored through a coverage memory $\boldsymbol{C}_k$ updated online from the executed trajectory, which is used to trigger replanning and to guide subsequent direction selections. As a result, inspection is not prescribed as a fixed viewpoint sequence, but adapts online to what has already been observed.

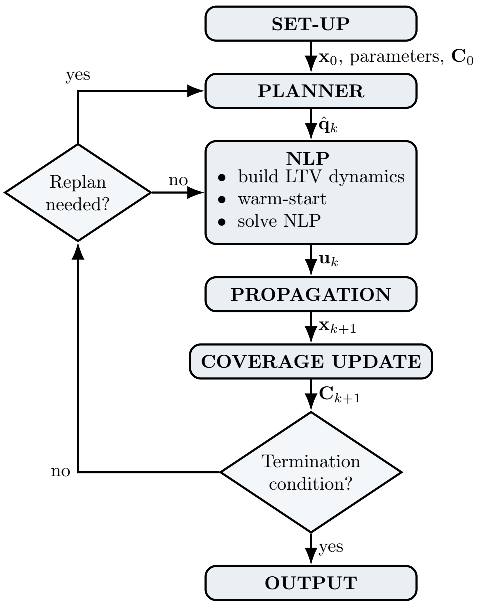

Figure 1: Closed-loop workflow.

Mapping strategy

Mapping progress is monitored on a viewing sphere by discretizing $\mathbb{S}^2$ into N_d unit directions $\{\boldsymbol{q}_i\}_{i=1}^{N_d}$. A spherical Fibonacci lattice generates a near-uniform distribution of directions, providing a simple baseline when the target geometry is not available a priori. A spherical Voronoi tessellation then assigns an angular region to each direction, turning the discrete set into a partition of the viewing domain into N_d mapping zones. To support proximity inspection, each angular zone is extended over an admissible standoff interval $[r_{\min}, r_{\max}]$, yielding volumetric sectors. These limits are set by imaging requirements to maintain a suitable target image scale.

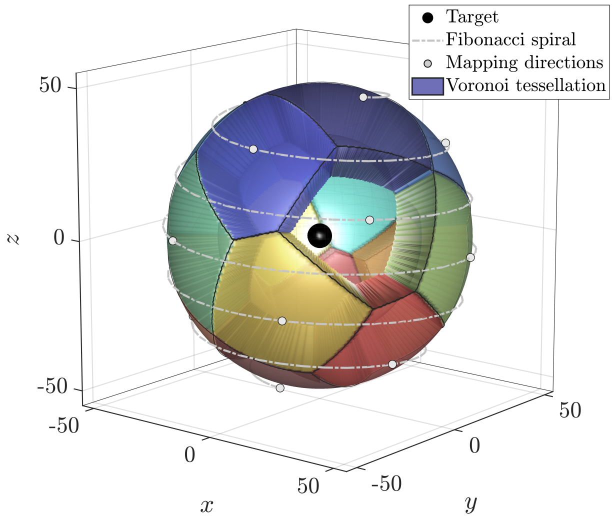

Figure 2: Imaging space discretization.

The online coverage memory is stored as a vector $\boldsymbol{C}_k \in \mathbb{R}^{N_d}$, whose i-th component accumulates the effective imaging time assigned to zone i. Coverage is incremented only when the current sample is usable for reconstruction, as assessed by an imaging gate $g_k \in [0, 1]$ built from range compliance within the standoff corridor and pointing compliance within the camera FoV. A smooth zone-membership vector $\boldsymbol{A}_k \in [0, 1]^{N_d}$ maps the current viewing direction to the Voronoi zones while avoiding hard switching at region boundaries. Accordingly, the per-step coverage increment is defined as

$$\Delta \boldsymbol{C}_k = g_k \, \Delta t \, \boldsymbol{A}_k, \qquad \boldsymbol{C}_{k+1} = \boldsymbol{C}_k + \Delta \boldsymbol{C}_k,$$

so that $\Delta \boldsymbol{C}_k \geq \mathbf{0}$ and coverage accumulates only under admissible imaging conditions. Progress is monitored by comparing $\boldsymbol{C}_k$ to a target per-zone level C_{tgt}, which provides a stable reference to prioritize under covered regions during replanning and to declare completion once all zones reach the required coverage.

High-level planner

The planner uses the current state estimate and the coverage memory to select the next exploration direction $\hat{\boldsymbol{q}}_k$ on the viewing sphere. Rather than computing a full inspection trajectory, it provides a compact, updatable reference that steers the guidance layer toward under covered regions while limiting manoeuvring. To meet runtime requirements, direction selection follows a coarse-to-fine hierarchy that increases model fidelity over a progressively reduced can-

didate set.
All candidate directions $\{\boldsymbol{q}_i\}$ are first screened using inexpensive proxies. The score combines a coverage deficit term, which estimates the mapping gain achievable along $\boldsymbol{q}_i$, with continuity penalties that discourage sharp changes with respect to the current motion direction, thus limiting excessive actuation. This stage outputs a shortlist of the top N_1 candidates.
The screened directions are then extended with a short two step lookahead. Sequences $(\boldsymbol{q}_0, \boldsymbol{q}_1)$ are formed from the shortlist by reapplying the same proxy score to the second step. For each first direction $\boldsymbol{q}_0$, all admissible follow-up candidates $\boldsymbol{q}_1$ are scored from the predicted end-of-leg state after committing to $\boldsymbol{q}_0$, using an updated heading for the continuity penalty and an updated coverage estimate.
Finally, the best N_2 sequences are refined through a short horizon constrained optimization. The sequence provides the viewing reference, while imaging requirements are assumed to be always satisfied. The refined score combines the maximum expected increment of $\boldsymbol{C}_{\text{pred}}$ under usable image assumptions with the control effort required by the optimized 3-DoF motion, and selects the best trade-off to define the next direction $\hat{\boldsymbol{q}}_k$.

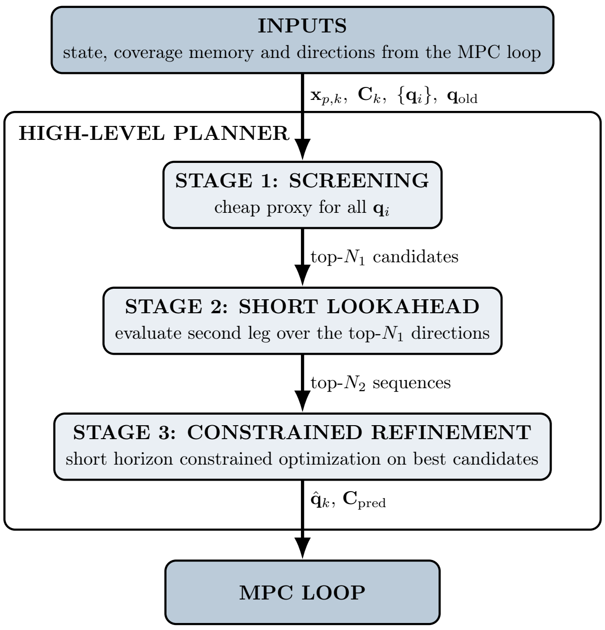

Figure 3: High-level planner architecture and MPC interface.

6-DoF Model Predictive Control

Given the exploration direction $\hat{\boldsymbol{q}}_k$, guidance is handled by a 6-DoF receding-horizon MPC. At each sampling time, a finite-horizon NLP is solved, only the first command is applied and the procedure is repeated in closed loop to track the planner selected direction while enforcing safety, actuation limits and imaging requirements.
The discrete time state and input are defined as

$$\boldsymbol{x}_k = \begin{bmatrix} \boldsymbol{\rho}_k^\top & \boldsymbol{v}_k^\top & \boldsymbol{\omega}_k^\top & \boldsymbol{\sigma}_k^\top \end{bmatrix}^\top, \quad \boldsymbol{u}_k = \begin{bmatrix} \boldsymbol{a}_k^\top & \boldsymbol{\tau}_k^\top \end{bmatrix}^\top$$

where $\boldsymbol{\rho}_k$ and $\boldsymbol{v}_k$ denote the relative position and velocity in the target centered LVLH frame, while $\boldsymbol{\omega}_k$ and $\boldsymbol{\sigma}_k$ represent the rotational state via body rates and MRPs. Translation is predicted using a discrete time LTI Clohessy–Wiltshire–Hill model, whereas attitude is predicted through a first order affine LTV approximation linearized along a reference rotational trajectory and frozen within each solve.
Imaging requirements are treated as usability conditions rather than hard constraints. They define the usable image regime adopted throughout the architecture, including coverage updates, and combine range regulation within $[r_{\min}, r_{\max}]$ shaped toward r_{opt} with boresight-to-LOS alignment consistent with the camera field of view. These thresholds follow from camera geometry and reconstruction needs and promote stable framing and pose consistency. The MPC cost therefore includes smooth terms promoting range keeping and pointing, a direction tracking term that drives the chaser toward the planner selected viewing direction $\hat{\boldsymbol{q}}_k$, and quadratic effort regularization on $\boldsymbol{a}_k$ and $\boldsymbol{\tau}_k$ to account for actuation usage and fuel consumption.
Safety and actuation envelopes are enforced through hard constraints together with discrete time prediction consistency. The previous solution is shifted and reused as a warm start, improving robustness and reducing solve time across consecutive iterations.
The applied command is propagated with a validation simulator integrating coupled nonlinear translational and rotational dynamics and including perturbations beyond the predictive model. This step provides the next state for the MPC cycle and assesses robustness under unmodeled effects, while driving the coverage up-

date through the real relative motion and attitude.

4. Results

The following results refer to a representative LEO inspection scenario. The non-cooperative target follows a near-circular Sun-synchronous orbit at altitude $h = 550$ km and inclination $i_t = 97.5°$, which defines the LVLH frame used for relative motion. Proximity operations are conducted within a prescribed standoff corridor and an inner keep-out radius, while mapping progress is monitored over $N_d = 16$ discretized viewing directions with a required per-direction coverage level $C_{\text{tgt}} = 10$ min, selected to provide sufficient imaging time and viewpoint diversity within each zone for reconstruction oriented data collection. The 6-DoF MPC runs at $T_s = 5$ s with prediction horizon $N = 200$. The sampling time is short enough for node-wise constraint enforcement to remain representative, while the horizon is long enough to capture the inspection dynamics.

Guidance performance

The chaser motion remains consistent with the intended close-proximity regime. The relative trajectory covers a wide angular region around the target, enabling viewpoint diversity and completes the mapping campaign in approximately 4.36 h. The out-of-plane component stays contained, indicating that cross-track excursions are used to enrich viewing geometry without aggressive repositioning. Relative velocity and body rate profiles remain within bounds and evolve smoothly, supporting imaging by limiting rapid translational drift and fast attitude motion. Control inputs respect saturation and rate limits, transients are localized around replanning events and do not lead to sustained saturation. Overall propellant usage is compact, with a total $\Delta v \simeq 1.81\ \mathrm{m\,s^{-1}}$ over the campaign.

Imaging behaviour

Imaging admissibility is assessed through range regulation and boresight-to-LOS alignment, consistently with the usable image regime adopted by the coverage update.
After a short transient, the relative range remains within the prescribed corridor and stays

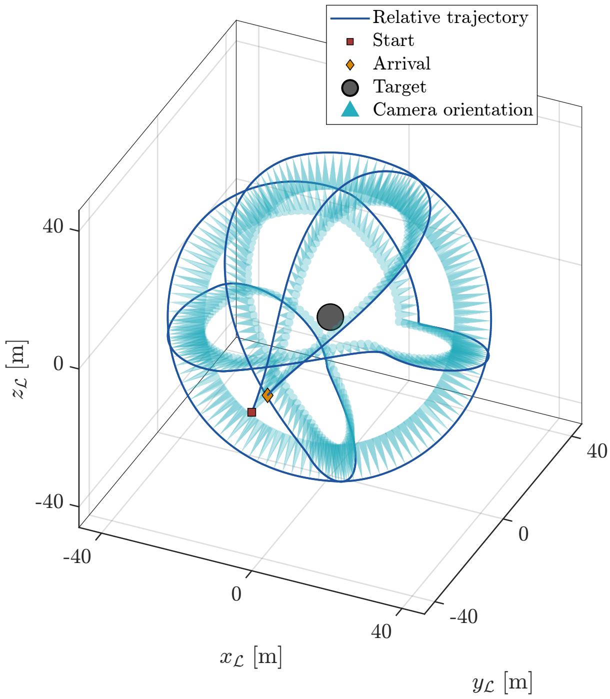

Figure 4: Relative trajectory in the target LVLH frame.

close to the nominal standoff distance r_{opt} over the campaign. Since range is promoted through objective shaping rather than imposed as a hard constraint, small deviations are allowed to accommodate viewpoint transitions and reorientation without forcing unnecessary control effort into strict range tracking.
Pointing performance is evaluated through the off-boresight angle between the camera boresight and the LOS. Alignment is maintained tightly for most of the run, keeping the off-boresight angle close to zero during steady inspection phases. Brief excursions occur mainly around direction updates and local geometric transitions, and remain bounded. These range and pointing trends indicate a sustained admissible imaging regime, so that coverage accumulation is driven primarily by usable image time.

Mapping progress and final coverage

Mapping progress is tracked through the coverage memory, which stores the accumulated usable image time associated with each discretized viewing direction. Completion is evaluated with a cap at the target level C_{tgt}, so that additional imaging time on already saturated directions does not inflate the progress measure. A scalar

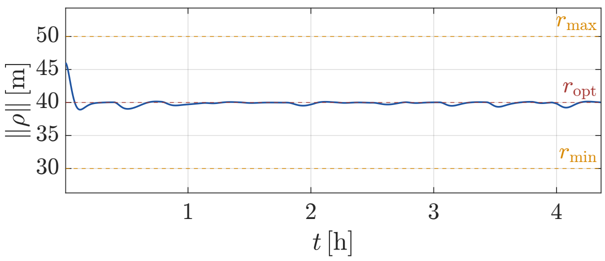

Figure 5: Relative range $r_k = \|\boldsymbol{\rho}_k\|$ over time with the admissible corridor and nominal distance.

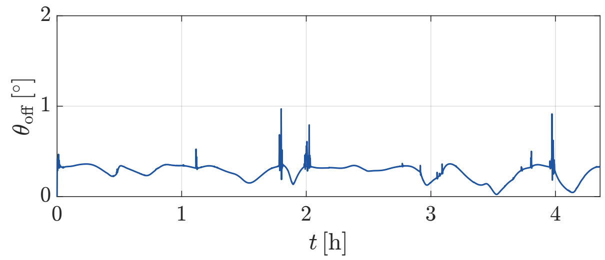

Figure 6: Off-boresight angle over time.

completion score is then obtained as the average capped coverage across all directions, and the mapping campaign terminates when this score reaches the prescribed acceptance level.

In the reported run, the completion score increases steadily and reaches the desired level within a finite duration. Progress is initially close to linear, since most directions are below C_{tgt} and admissible imaging conditions are readily available. As directions saturate, the increase becomes more intermittent, with brief plateaus caused by continuity preserving transitions that temporarily revisit covered sectors. The capped score advances only when time is accrued on directions that have not yet reached C_{tgt}.

The final outcome is summarized by the capped per-direction distribution in Figure 7. The stopping rule guarantees the prescribed global completion level at termination and, in this simulation, all directions reach at least C_{tgt}. This is supported by the return leg toward the initial viewpoint, which contributes usable imaging time and helps close residual deficits. The remaining dispersion is expected in closed loop and reflects unavoidable revisits and continuity requirements, since perfectly equal per-direction values would require disproportionate reconfiguration effort.

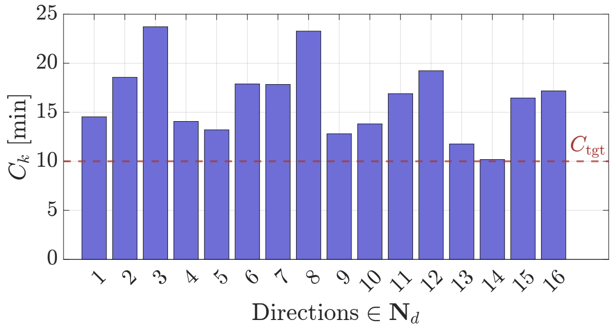

Figure 7: Final coverage distribution $\boldsymbol{C}_k$ across discretized viewing directions.

Planner behaviour

In closed loop, the direction is updated only when progress-based conditions are met, so replanning acts as an episodic decision process rather than as a high-frequency correction. As a result, direction changes are sparse with respect to the MPC sampling rate and rapid back-and-forth switching between neighbouring zones is avoided.

As expected, instead of always selecting the most under covered region at each trigger, the planner trades coverage improvement against the feasibility and cost of the transition in close-proximity. This yields an approximately cyclic progression across complementary sectors, where viewpoint continuity is preserved and reconfigurations remain limited. Overall, the observed behaviour is consistent with a policy that promotes coverage completion while containing manoeuvring effort by exploiting the natural evolution of the relative geometry.

Online feasibility

Since the MPC policy is applied at the sampling rate T_s, real-time operation requires that each per-step solve completes within the available time budget, while event-driven replanning remains compatible with the same constraint. All computational performances are reported for a fixed solver configuration.

At the MPC level, solve times remain consistently below the sampling interval. Occasional peaks arise from localized changes in constraint activity and reference updates, yet they do not challenge the timing requirement even when reserving margin for state estimation and command uplink within the same cycle.

The high-level planner adds overhead only when a direction update is triggered. In the present implementation, a replanning call typically requires on the order of 15–25 s, but the updated direction is treated as immediately available at the next MPC solve for the results reported here. In an onboard setting, the same latency can be handled without altering the formulation by running the planner asynchronously over upcoming MPC intervals or by holding the current direction until an update is available.

3D reconstruction

To validate the proposed imaging driven guidance, the acquired frames are assembled into a dataset and processed with a NeRF-based reconstruction pipeline. The resulting model remains visually consistent across representative views, supporting that the trajectory provides sufficiently diverse and well distributed imagery.

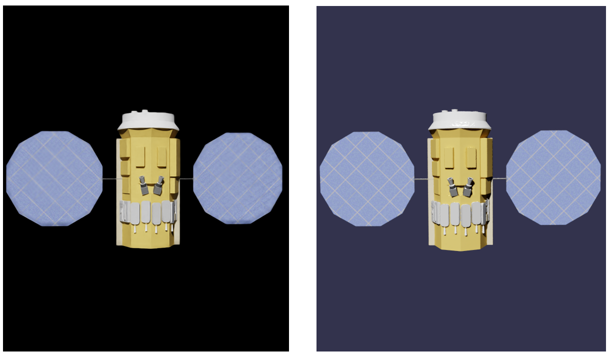

(a) Dataset frame. (b) Mesh snapshot.

Figure 8: Qualitative NeRF-based validation: a sampled dataset image (a) and a CAD snapshot of the reconstructed 3D mesh (b).

5. Conclusions

This thesis addressed autonomous close-proximity inspection around an unknown, non-cooperative space object, with the operational goal of generating a viewpoint set suitable for deep learning based 3D reconstruction. The core contribution is integrating inspection objectives and constrained guidance within a single closed-loop architecture, so that feasible trajectories also deliver informative viewpoints. Numerical results confirm bounded motion and actuation, sustained imaging admissibility, and completion of the prescribed mapping acceptance condition within finite time. Crucially, the approach supports autonomy under partial prior knowledge, since exploration directions are generated online from coverage information and the same hierarchy can incorporate refined estimates of safety and reconstruction parameters as inspection knowledge evolves.

Future developments mainly target increased flight realism. Planner latency motivates faster screening and refinement, together with asynchronous execution under a fixed deadline and conservative fallback when an update is delayed. The imaging layer can also be extended beyond geometric proxies by accounting for illumination and visibility and by introducing uncertainty aware coupling between perception and guidance as geometry estimates improve.

Overall, this work provides an onboard oriented baseline for imaging driven close-proximity inspection and mapping, combining a persistent representation of mapping progress with constrained receding-horizon control. The resulting architecture enables a principled trade between mapping completeness, safety margins, and actuation effort, and can be adapted to a broad set of inspection and characterization objectives beyond the specific reconstruction pipeline considered.

References

[1] Bala Prenith Reddy Gopu, Timothy J Huber, George M Nehma, Patrick D Quinn, Madhur Tiwari, Matt Ueckermann, David Hinckley, and Christopher McKenna. Dynamic scene 3d reconstruction of an uncooperative resident space object. In *AIAA SCITECH 2026 Forum*, page 1038, 2026.

[2] Michele Maestrini and Pierluigi Di Lizia. Guidance strategy for autonomous inspection of unknown non-cooperative resident space objects. *Journal of Guidance, Control, and Dynamics*, 45(6):1112–1127, 2022.

[3] Antonio Rizza, Francesco Topputo, and Simone D'Amico. Goal-oriented asteroid mapping under uncertainties using sequential convex programming, 01 2024.

[4] Ding Zhou, Yanquan Zhang, and Shunli Li. Receding horizon guidance and control using sequential convex programming for spacecraft 6-dof close proximity. *Aerospace Science and Technology*, 87:459–477, 2019.