

An Adaptive Inspection Planning Approach Towards Routine Monitoring in Uncertain Environments

Vignesh Kottayam Viswanathan, Yifan Bai, Scott Fredriksson, Sumeet Satpute,
Christoforos Kanellakis and George Nikolakopoulos

Abstract—In this work, we present a hierarchical framework designed to support robotic inspection under environment uncertainty. By leveraging a known environment model, existing methods plan and safely track inspection routes to visit points of interest. However, discrepancies between the model and actual site conditions, caused by either natural or human activities, can alter the surface morphology or introduce path obstructions. To address this challenge, the proposed framework divides the inspection task into: (a) generating the initial global view-plan for region of interests based on a historical map and (b) local view replanning to adapt to the current morphology of the inspection scene. The proposed hierarchy preserves global coverage objectives while enabling reactive adaptation to the local surface morphology. This enables the local autonomy to remain robust against environment uncertainty and complete the inspection tasks. We validate the approach through deployments in real-world subterranean mines using quadrupedal robot. A supplementary media highlighting the proposed method can be found here https://youtu.be/6TxK8S_83Lw.

I. INTRODUCTION

Robotic inspection plays a critical role in modern mining operations, with robots increasingly deployed to handle unsafe and repetitive tasks: from detecting changes in subterranean environments [1], [2] to routine monitoring [3], [4]. This work contributes to enabling autonomous visual inspection in uncertain subterranean environments, focusing specifically on surface inspection and the discrepancies that arise between a priori models and current environmental conditions. These discrepancies may result from obstructions along pre-recorded paths (e.g., rockfalls) or from routine mining activities that alter surface morphology (e.g., ore extraction). Since such uncertainty is often unpredictable and cannot be fully modelled in advance, it can degrade inspection performance and compromise the quality of visual data collected during missions. Figure 1 presents a few demonstrative use-case of robotic inspection application in subterranean mines.

Therefore, addressing visual inspection tasks in a potentially outdated environment model poses two main challenges. First, discrepancies between current conditions and historical environment data can compromise the generated inspection route. These mismatches may introduce obstacles that either obstruct the path or block sensor views. As a result, the system must replan locally to ensure coverage of

The authors are with Robotics and AI, Luleå University of Technology 97187, Luleå, Sweden {vigkot, yifan.bai, scofre, sumsat, chrkan and geonik}@ltu.se

The project is funded by the European Union’s Horizon Europe Research and Innovation Programme, under the Grant Agreement No. 101138451 PERSEPHONE

the desired region in a safe and reliable manner. Second, changes in the surface morphology can directly affect the resolution of visual data collected at each view-pose. The replanned route relies on viewing requirements derived from the outdated environment model. When actual surface conditions differ from the model, the inspection plan must be updated to reflect the current environment conditions.

Collectively, such morphological discrepancies translates into actions necessitating multiple mapping runs: to update the model, plan on the latest representation and finally execute the plan in the environment. As such, they contribute to an overall longer mission duration. To address these challenges, we propose a hierarchical inspection framework that aims to address the problem of routine inspection in uncertain environments. This is achieved by leveraging a priori 3D map information for the initial task planning and implementing a local reactive view planner for online replanning. Through the proposed methodology, we aim to actively address the model mismatch problem through an adaptive planning approach within a single session.

The article is structured as follows: Section II presents the review on state-of-art methods and the main contributions of the article. Section III highlights the problem statement and the proposed methodology in detail. Section IV and V presents the evaluation setup, the experimental results and the related discussions. Finally, Sec VI underscores the insights drawn from the experiments and highlights the future works of the current research.



Fig. 1. Examples of robotic inspection application in subterranean environments.

II. RELATED WORKS

[5] presents a hierarchical inspection framework designed to provide coverage of target regions in a 3D scene. In their work, the authors propose an initial viewpoint sampling covering the desired inspection regions, followed by clustering the valid collision-free viewpoints within each environment subspace representation. Subsequently, a Travelling Salesman Problem (TSP) tour is solved and later refined to provide shortest collision-free path to cover the desired targets. Similar to this, [6] propose a model based hierarchical approach by decomposing the 3D scene for efficient viewpoint sampling and subsequent path planning for complete coverage. Specifically, the authors leverage an input 3D point-cloud of the target scene which is then skeletonized to sample viewpoints. This is then processed to extract and execute the shortest collision-free path covering the scene. In [7], the authors present an uncertainty aware inspection framework targeted towards adaptive view planning strategy in presence of static and dynamic obstacles in the scene. Particularly, the authors leverage a priori map to compute a global view plan of the scene and then adapt it online to ensure safe navigation and occlusion-free viewing during the mission.

While [5]–[7] demonstrate autonomous coverage of a 3D scene based on an a priori model, limited works [7] address the presence of uncertainty in environments. However, these approaches often overlook a broader aspect of environmental uncertainty, namely, evolving surface morphologies. In contrast, our approach uses an a priori model to define inspection tasks and generate a global view plan. This is integrated with a reactive local view planner to maintain awareness of surface morphology and adapt in real-time, ensuring desired photogrammetric requirements are met during coverage of the inspection region.

A. Contributions

Building on the above discussion, we outline our key contributions that address the challenges of autonomous inspection in uncertain environments. For clarity, we define the environmental uncertainty as the incomplete or outdated information (e.g., lack of presence of obstacles, absence of knowledge on current surface morphology) available to the robot or human operator during inspection planning.

To address the above limitations, we propose a hierarchical inspection framework that combines global planning based on a historical model with local online replanning to adapt to unexpected environmental changes. Given a priori map and target inspection regions, we generate corresponding global view plans before mission start. Each inspection task is then passed to a local planner, which generates a predicted inspection route using real-time sensor data. During each task, we compute the similarity between the predicted and replanned inspection paths, quantifying the deviation observed. When significant deviation is perceived, the system prioritizes local replanning to update the preplanned path to preserve the desired photogrammetric constraints. Throughout replanning, the system ensures awareness of the original global plan and only executes replanned segments evaluated to deviate based on the current perception of the environments. Holistically, this approach enables adaptive surface inspection in the presence of unexpected environmental changes.

Additionally, we demonstrate the efficacy of the proposed framework in real-world applications. We achieve this through field deployment experiments on a quadrupedal robot in subterranean environments. Furthermore, we evaluate the framework in large-scale missions in simulated scenarios.

III. METHODOLOGY

A. Preliminaries

We denote the robot's pose estimate as $\hat{\mathbf{X}}_{odom}^k = [x, y, z, \phi, \theta, \psi]^T \in \mathbb{R}^6$. The camera horizontal and vertical Field-Of-View (FOV) are $\alpha, \beta \in \mathbb{R}_+$, with the desired viewing distance $d_{view} \in \mathbb{R}_+$ and image overlaps $\gamma_H, \gamma_V \in \mathbb{R}_+$. The input point clouds are $\mathbf{p}^k = \{\mathbf{p}_i^k \in \mathbb{R}^3\}_{i=1}^M$. A reference view-pose is $\mathbf{X}_{ref} \in \mathbb{R}^4 | \mathbf{X}_{ref} = [x, y, z, \psi]$, and the predicted local path is $\Pi_{LVP}^k = \{\mathbf{X}_{ref}^{k+1}, \dots, \mathbf{X}_{ref}^{k+N}\}$, where $N \in \mathbb{Z}_+$ is the prediction horizon.

B. Problem Statement

For a known, historical volumetric map of the environment $\mathbb{M}_H \in \mathbb{R}^3$, we denote $\mathbf{R} = \{\mathbf{v}_i\}_{i=1}^M | \mathbf{v}_i \in \mathbb{R}^3$ as a simple flat 3D polygon representing the inspection region of interest in \mathbb{M}_H . Given a view-pose as $\mathbf{x} = [x, y, z, \psi] \in \mathbb{R}^4$, we define \mathbf{Z} as the initial set of view poses generated for \mathbf{R} over \mathbb{M}_H . We denote $\mathbb{M}^\delta \in \mathbb{R}^3$ as the uncertainty brought about by

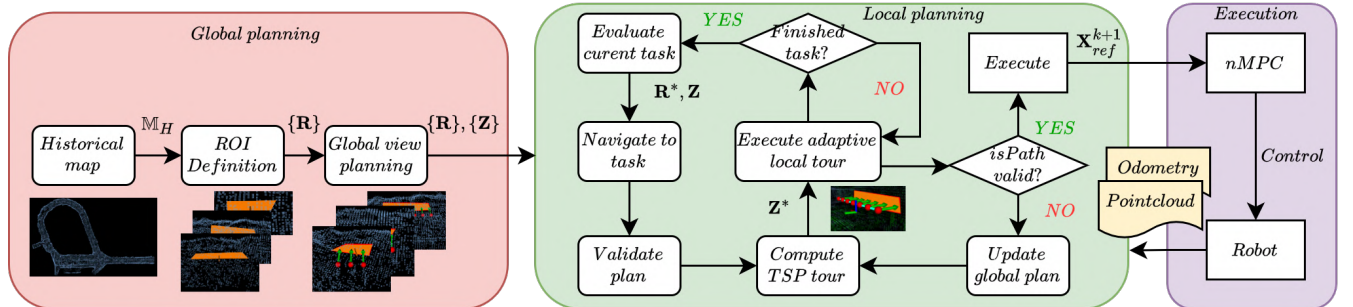


Fig. 2. Framework of the proposed methodology.

unexpected changes in the current scene. \mathbb{M}^δ comprises of changes from the previously known scene, accounting for path obstructions or evolution of the mine-face topology. We denote the current 3D representation of the environment, built by the robot as it navigates the scene, as $\mathbb{M}_C \in \mathbb{R}^3$. Thus, by defining an inspection task over \mathbb{M}_H , the presence of \mathbb{M}^δ can influence the validity of \mathbf{Z} . The goal of this work is to compute and execute an updated view plan respecting the viewing constraints in an online manner, evaluated to be necessary in presence of \mathbb{M}^δ and adapted to current environment conditions. Figure 2 presents the framework of the proposed methodology to address the problem of inspection in uncertain environments.

C. Hierarchical Inspection Planning

Global View-Planning: Figure 3 presents the implemented initial view plan generation scheme based on [8]. Subsequent to the initial ROI definition \mathbf{R} (Fig 3(a)), the region is discretized into grids (Fig 3(b)). We consider the camera footprint and the desired photogrammetric constraints to generate the vertical and horizontal grid lines. The distance between the grid lines is modelled based on the required overlap distance between two successive view poses. Next, we filter the grid intersections to keep only those that lie within the polygon region (Fig 3(c, left)). The valid intersections are then used to compute the corresponding set of view poses \mathbf{Z} . To achieve this, each valid grid intersection is projected along the direction of the normal vector of the polygon. The poses are projected to the desired viewing distance and oriented to view the ROI surface (Fig 3(c, right)).

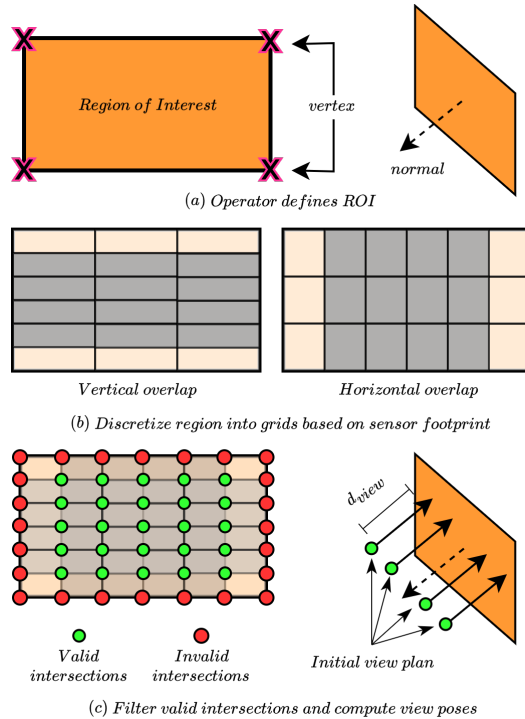


Fig. 3. A visual representation of the initial view plan generation scheme for a defined region of interest.

Local View Planning: We leverage the reactive local view planner initially established in [9]. In brief, the planner operates over instantaneous 3D point cloud measurements to generate a reference view pose satisfying the desired viewing constraints. Specifically, at each planning instance, \mathbf{p}^k is processed to evaluate the nearest 3D point on the locally observed surface $\mathbf{p}_{nn}^k \subseteq \mathbf{p}^k$. From this, the reference unit vectors capturing the ego orientation of the robot is determined. Based on the viewing constraints, $d_{view}, \gamma_H, \gamma_V$, the next view-pose in sequence is computed with respect to the current position of the robot. Equation (1) presents the mathematical representation of the view planner. For a detailed insight on the local view planner, the readers are directed to [9].

$$\mathbf{X}_{ref[x,y,z]}^{k+1} = \hat{\mathbf{X}}_{odom[x,y,z]}^k + \vec{\nu}_x d_{insp} + \vec{\nu}_y d_{hov} + \vec{\nu}_z d_{vov} \quad (1)$$

where,

$$\begin{aligned} \vec{\nu}_x^k &= \frac{\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k}{\|\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k\|} \\ \vec{\nu}_y^k &= \vec{\nu}_{up} \times \vec{\nu}_x^k \\ \vec{\nu}_z^k &= \vec{\nu}_x^k \times \vec{\nu}_y^k \\ d_{hov} &= 2 \tan\left(\frac{\alpha}{2}\right) \|\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k\| \\ &\quad - 2 \tan\left(\frac{\alpha}{2}\right) \|\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k\| \gamma_H \\ d_{vov} &= 2 \tan\left(\frac{\beta}{2}\right) \|\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k\| \\ &\quad - 2 \tan\left(\frac{\beta}{2}\right) \|\mathbf{p}_{nn}^k - \hat{\mathbf{X}}_{odom[x,y,z]}^k\| \gamma_V \end{aligned}$$

such that, $\vec{\nu}_x, \vec{\nu}_y, \vec{\nu}_z \in \mathbb{R}^3$ are the unit view vectors of the robot with $\vec{\nu}_{up} = [0, 0, 1]$ being the unit pointing vector along the +Z-axis of the platform. Equation (2) presents the required yaw orientation to be maintained by the robot during local planning.

$$\mathbf{X}_{ref[\psi]}^{k+1} = \arctan(\vec{\nu}_x^{k+1}(1), \vec{\nu}_x^{k+1}(0)) \quad (2)$$

Planning under uncertainty: Figure 4 presents the pipeline for adaptive local planning in presence of environment uncertainty. Once the tasks are defined and the mission initialized, we evaluate and compute a global traversable pathway to the tasks from the current position over \mathbb{M}_H (refer Fig 4(b)). The traversable route length is used to prioritize the closest task \mathbf{R}^* respective to the robot's position. To ensure safe navigation in the scene, we segment the global path and compute and execute the new traversable route over \mathbb{M}_C . To generate a pathways graph, we adapt the implementation presented in [10]. If an inspection task is evaluated to be unreachable or within collision region, we update the task order and proceed to the next valid task. Though this work uses total traversable distance as the primary weighting metric to rank candidate tasks, the framework is objective-agnostic and can incorporate alternative or additional criteria to improve decision-making.




Fig. 4. A visual schema of the proposed adaptive inspection planning approach in uncertain environments.

Once at target position, the effective visiting tour Z^* , is computed from Z . This is solved as a Simulated Annealing Travelling Salesman Problem (SA-TSP) [11]. As the robot starts the tour, we extract a corresponding subset of the global inspection path $\Pi_{GVP}^k = \{x_k^*\}_{k=1}^{k+N} | x_k^* \in Z^*$. In order to predict the local inspection path Π_{LVP}^k , we utilize Π_{GVP}^k to guide the view plan generation in (1). The predicted inspection path is then computed by recursively evaluating (1) with the updated position over the prediction horizon N . We then compare this segment against the locally planned path Π_{LVP}^k to assess the validity of Z (refer Fig 4(c)). Since Π_{LVP}^k is derived from instantaneous point-cloud measurements (see (1)), morphological changes in the inspection surface directly influence its geometry. As a result, Π_{LVP}^k adapts to the current environment and deviates geometrically from Π_{GVP}^k .

We model this deviation by computing the Fréchet distance ($F_D \in \mathbb{R}_+$) between the two paths. We integrate this measure of similarity as $\Gamma_s \in \mathbb{R}_+ | \Gamma_s = [0, 1]$ as defined in (3).

$$\Gamma_s = \frac{1}{1 + F_D(\Pi_{GVP}^k, \Pi_{LVP}^k)} \quad (3)$$

where, F_D is modelled as presented in [12].

Higher values of Γ_s indicate greater deviation between the global and local paths, implying increased environmental uncertainty compared to the prior map \mathbb{M}_H . To adapt planning in the presence of such uncertainty, we define a threshold $\Gamma_t \in \mathbb{R}_+$, above which the local path Π_{LVP}^k is considered as the candidate path to be tracked for the current inspection task.

Although Equation (3) does not quantify uncertainty explicitly, the similarity metric Γ_s offers an indirect measure of discrepancy in the environment, as path deviations reflect local changes in surface morphology. This introduces the challenge of keeping the inspection effort confined to the originally defined inspection region R^* . To reconcile the updated local plan Π_{LVP}^k with this region, we leverage an approximated global path $\hat{\Pi}_{GVP}^k$ to guide the replanning effort via the Kabsch method [13] (refer Fig 4(d)). Equation (4) presents the optimal alignment used to reproject Π_{GVP}^k in case where $\Gamma_s > \Gamma_t$.

$$\hat{\Pi}_{GVP}^k = \begin{cases} U^* \Pi_{GVP}^k + t^* & , \text{if } \Gamma_s > \Gamma_t, \\ \Pi_{GVP}^k & , \text{otherwise} \end{cases} \quad (4)$$

where, the optimal rotation matrix $U^* \in SO(3)$ and the translation vector $t^* \in \mathbb{R}^3$ are the rigid transform obtained from the Kabsch method applied to align Π_{GVP}^k to Π_{LVP}^k .

Finally, the reference view plan is fed to a nonlinear model Predictive Controller (nMPC) [14], where the first view pose is tracked and the process (1)-(4) repeats until the inspection task is completed. We consider the completion criteria as the robot tracking the global visitation tour Z^* , either directly or through approximation. The pseudocode of the above process is provided in Alg 1.

Algorithm 1: Adaptive Surface Inspection

Input: $\mathbb{M}_H, \alpha, \beta, \gamma_H, \gamma_V, d_{view}$
 $\{R\} \leftarrow \text{DefineInspectionTask}(\mathbb{M}_H)$
 $\{Z\} \leftarrow \text{GenerateGlobalViewPlan}(\{R\})$
 $R^*, Z \leftarrow \text{EvalInspectionTasks}(\{R\}, \{Z\})$
NavigateToTask(R^*)
if *not* isCollisionFree(Z) **then**
 | **UpdatePlan**(Z)
 $Z^* \leftarrow \text{ComputeVisitingTour}(Z)$
 $\Pi_{GVP}^k \leftarrow \text{ExtractGlobalSegment}(Z^*)$
 $\Pi_{LVP}^k \leftarrow \text{PredictLocalPlan}(\hat{X}_{odom}^k, P^k, \Pi_{GVP}^k)$
 $\Gamma_s \leftarrow \text{ComputePathSimilarity}(\Pi_{GVP}^k, \Pi_{LVP}^k)$
if $\Gamma_s > \Gamma_t$ **then**
 | $\hat{\Pi}_{GVP}^k \leftarrow \text{ComputeAlignment}(\Pi_{GVP}^k, \Pi_{LVP}^k)$
 | $nMPC \leftarrow \text{ExecuteTracking}(\Pi_{LVP}^k)$
 | **CheckThreshold**($\hat{\Pi}_{GVP}^k$) **youT**
else
 | $nMPC \leftarrow \text{ExecuteTracking}(\Pi_{GVP}^k)$
 | **CheckThreshold**(Π_{GVP}^k)
 $k + 1 \leftarrow k$

IV. SETUP AND EVALUATION

The evaluation of the proposed methodology is carried out in real-world underground mines (refer Fig 5). We evaluate along two fronts: (i) task-level inspection performance with the proposed hierarchical planner, and (ii) adaptivity to static and dynamic environmental uncertainty (e.g., surface evolution and newly introduced obstacles). To measure the quality of observation from a viewpoint, we consider the angle of incidence subtended per-pixel normal to the camera view direction. We follow the method similar to the one presented in [15] to estimate the surface normal from the




Fig. 5. A visualization of the subterranean environment.

depth image. Subsequently, the viewpoint utility is computed as the mean of the cosine of the surface incidence angle, computed per valid pixel from the depth image. Throughout the runs, we consider $d_{view} = 2\text{m}$, $\gamma_H = 0.6$, $\alpha = 69.5^\circ$, $\beta = 45^\circ$ and $\Gamma_t = 0.5$.

The architecture is evaluated onboard Boston Dynamics (BD) Spot quadruped robot equipped with Orbbec Gemini2XL stereo-camera, Vectornav VN-100 Inertial Measurement Unit (IMU) and Ouster OS-0 3D LiDAR. We use Intel NUC embedded computational board running ROS Noetic and Ubuntu 20.04. We use Voxelblox [16] for 3D mapping. All planning tasks are evaluated on the onboard computational unit, from loading prior map to defining and evaluating tasks and finally executing inspection of the desired inspection regions.

V. RESULTS AND DISCUSSIONS

Figure 6 presents a collage of the steps taken by the hierarchical framework during deployment onboard BD SPOT in real-world subterranean mine. Fig 6(a) presents the loading of the historical map \mathbb{M}_H from a previous run. The inspection tasks $\{\mathbf{R}\}$ are then defined within \mathbb{M}_H and the global view plan $\{\mathbf{Z}\}$ is computed (refer Fig 6(b)). Figure 6(c) visualizes the global traversable pathway computed for each task over

\mathbb{M}_H . Figure 6(d) visualizes these color-coded paths based on the total traversable distance from the robot's current position. Once the candidate task is evaluated, the planner proceeds to segment, recompute and execute an updated path to \mathbf{R}^* over \mathbb{M}_C .

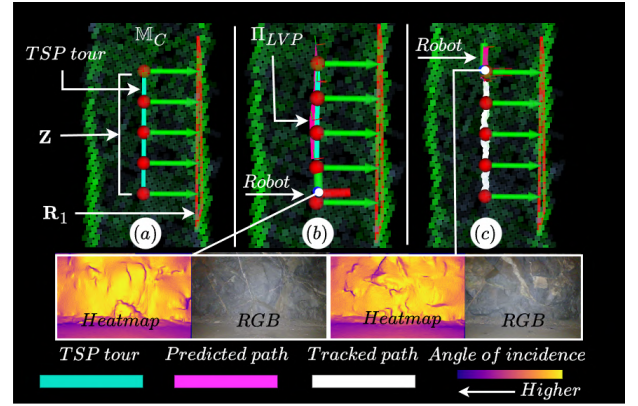


Fig. 7. A visual collage capturing the inspection behaviour for region \mathbf{R}_1 .

Figure 7 and 8 present a nominal scenario with inspection performance for the region \mathbf{R}_1 . The task was completed in 30 seconds with the planner evaluating Γ_s below the defined threshold value of 0.5 (refer Fig 8(middle)). As such, the planner followed the global plan \mathbf{Z}^* . This is reflected in the profile of the predicted local inspection path (refer Fig 7(b)) aligning with the global view plan during the inspection of \mathbf{R}_1 . The viewpoint utility, defined earlier in Sec IV, records an average value of 0.71, with values closer to 1 indicating better orientation of the camera sensor towards the local surface.

For the case of addressing inspection of evolving surface morphology, Fig 9 and 10 captures the adaptive nature of the proposed methodology. Figure 10(a) presents the historical map built in a previous run. To simulate evolving surface




Fig. 6. A collage highlighting the initial steps of the proposed hierarchical inspection framework during real-world deployment in a subterranean mine.

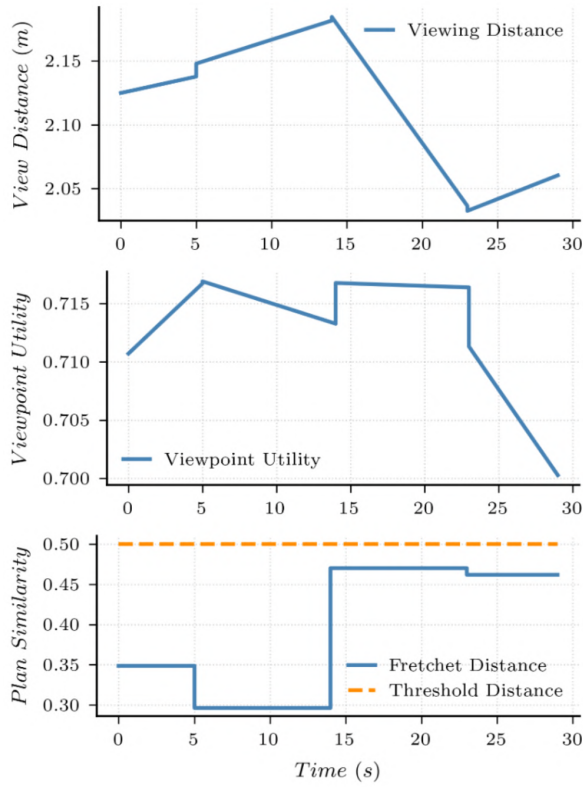


Fig. 8. Inspection performance recorded for R_1 .

morphology, we modify the scene to reflect a receding surface effect. The highlighted section in Fig 10(a) captures the introduced objects in the scene. Once \mathbb{M}_H is built, we remove the objects to simulate a receding surface. Figure 10(b) presents the inspection region definition and subsequent computation of the global view plan. Since only a single task is defined in this scenario, the planner computes the single global pathway to the task upon mission initialization (refer Fig 10(c)).

In terms of quantitative analysis, Figure 9 captures the performance metrics during inspection. Particularly, the adaptive behaviour is reflected in the quick maintenance of desired viewing distance as the robot starts tracking the global view plan due to the shift in the registered surface from \mathbb{M}_H . This is brought into effect due to the higher Γ_s score evaluated during inspection. From $t = 20s$, the Root Mean Square Error (RMSE) between the global view plan path Π_{GVP} and the predicted local path Π_{LVP} before Kabsch alignment and the RMSE between $\hat{\Pi}_{GVP}$ and Π_{LVP} after Kabsch alignment begins to converge. This indicates that the robot is moving towards inspection of a surface present in both \mathbb{M}_C and \mathbb{M}_H . The small rise in the maintained viewing distance at 35s is due to the robot recapturing the global view plan since Γ_s drops below the defined threshold.

Figure 12 presents the adaptive inspection executed for the above scenario. In Fig 12(a), the robot's start position is highlighted alongside the main geometric difference between \mathbb{M}_H and \mathbb{M}_C . The removed objects leave behind previously unmapped regions in \mathbb{M}_H . From Fig 12(a)- 12(d), the adaptive nature of the proposed methodology is shown. Once the




Fig. 9. Inspection performance recorded for the scenario of inspection under evolving surface morphologies.

mission starts, the planner immediately replans an updated route to meet the desired viewing constraints. This process recurses until the estimated deviation between the predicted and global view plan drops below the defined threshold. Consequently, the robot starts tracking the global tour (refer Fig 12(e)).

We compare the performance of the proposed framework against a baseline method under environmental uncertainty to highlight the improvements in view quality obtained. For baseline, we consider the robot tracking the global tour directly without the adaptive replanning. Figure 11 captures the inspection behaviour without adaptive replanning. Contrasting against the results obtained in Fig 9, we can directly infer that the adaptive approach maintains a higher viewpoint utility score once the replanning starts. Furthermore, the proposed methodology rapidly attempts to close the offset in the viewing distance and drops below the threshold distance in 15s.




Fig. 10. A visual representation of the initial inspection setup demonstrated for modified surface morphology.

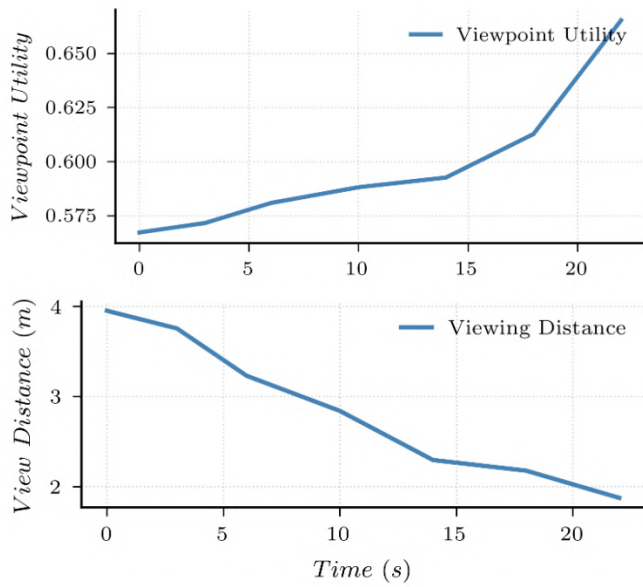


Fig. 11. Inspection performance metrics recorded for a non-adaptive inspection strategy under environment uncertainty.

VI. CONCLUSIONS AND FUTURE WORKS

In this work, we presented a novel approach for routine inspection in uncertain environments. Specifically, we demonstrated an adaptive inspection planning framework capable of responding to evolving surface morphologies. We showcase our work through real-world deployments. Our method integrates a hierarchical planning structure that combines global task definition and view planning with a reactive local layer, enabling effective adaptation to dynamic surface conditions. Through the evaluated scenarios, we demonstrate the benefits of an adaptive planning approach to address environment uncertainty in a single mission. We

also showcased the improvement in view quality achieved by our approach compared to a non-adaptive baseline.

For future work, we aim to extend the proposed methodology to multi-agent systems for large-scale inspection missions. Additionally, we plan to conduct more extensive field evaluations in diverse and unstructured environments to ensure robust, environment-agnostic performance.

REFERENCES

- [1] N. Stathouloupoulos, A. Koval, and G. Nikolakopoulos, "Irregular change detection in sparse bi-temporal point clouds using learned place recognition descriptors and point-to-voxel comparison," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8077–8083, IEEE, 2023.
- [2] L. Attard, C. J. Debono, G. Valentino, and M. Di Castro, "Vision-based change detection for inspection of tunnel liners," *Automation in Construction*, vol. 91, pp. 142–154, 2018.
- [3] M. Staniaszek, T. Flatscher, J. Rowell, H. Niu, W. Liu, Y. You, M. Gadd, M. Mattamala, A. Schutz, D. De Martini, *et al.*, "Autospect: Towards long-term autonomous inspection and monitoring," *IEEE Transactions on Field Robotics*, 2025.
- [4] S. Nordström, N. Stathouloupoulos, N. Dahlquist, B. Lindqvist, I. Tevetzidis, C. Kanellakis, and G. Nikolakopoulos, "Safety inspections and gas monitoring in hazardous mining areas shortly after blasting using autonomous uavs," *Journal of Field Robotics*, 2024.
- [5] C. Cao, J. Zhang, M. Travers, and H. Choset, "Hierarchical coverage path planning in complex 3d environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3206–3212, IEEE, 2020.
- [6] C. Feng, H. Li, M. Zhang, X. Chen, B. Zhou, and S. Shen, "Fc-planner: A skeleton-guided planning framework for fast aerial coverage of complex 3d scenes," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8686–8692, IEEE, 2024.
- [7] H. Jin, Z. Xu, H. Shen, X. Han, K. Ye, and K. Shimada, "Adaptive planning framework for uav-based surface inspection in partially unknown indoor environments," *arXiv preprint arXiv:2504.09294*, 2025.
- [8] S. S. Mansouri, C. Kanellakis, G. Georgoulas, D. Komiinak, T. Gustafsson, and G. Nikolakopoulos, "2d visual area coverage and path planning coupled with camera footprints," *Control Engineering Practice*, vol. 75, pp. 1–16, 2018.
- [9] V. K. Viswanathan, V. Sumathy, C. Kanellakis, and G. Nikolakopoulos, "A surface adaptive first-look inspection planner for autonomous




Fig. 12. A collage capturing the response from the adaptive inspection strategy demonstrated during field deployment.

- remote sensing of open-pit mines,” in *2024 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 280–285, 2024.
- [10] H. Oleynikova, Z. Taylor, R. Siegwart, and J. Nieto, “Safe local exploration for replanning in cluttered unknown environments for microaerial vehicles,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1474–1481, 2018.
- [11] Y. Bai, C. Kanellakis, and G. Nikolakopoulos, “Sa-recbs: Multi-robot task assignment with integrated reactive path generation,” *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 7032–7037, 2023.
- [12] T. Eiter, H. Mannila, *et al.*, “Computing discrete fréchet distance,” 1994.
- [13] W. Kabsch, “A solution for the best rotation to relate two sets of vectors,” *Foundations of Crystallography*, vol. 32, no. 5, pp. 922–923, 1976.
- [14] S. Karlsson, B. Lindqvist, and G. Nikolakopoulos, “Ensuring robot-human safety for the bd spot using active visual tracking and nmpc with velocity obstacles,” *IEEE Access*, vol. 10, pp. 100224–100233, 2022.
- [15] Y. Nakagawa, H. Uchiyama, H. Nagahara, and R.-I. Taniguchi, “Estimating surface normals with depth image gradients for fast and accurate registration,” in *2015 International Conference on 3D Vision*, pp. 640–647, IEEE, 2015.
- [16] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, “Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1366–1373, IEEE, 2017.