Variable-Resolution Virtual Map Guided Informed Tree Search for Autonomous USV Exploration

Ye Li[1], Mal Fazliu[2], Yewei Huang[3,4], Alberto Quattrini Li[3], Brendan Englot[4], Cunjia Liu[2], and Yuanchang Liu[1]

***Abstract*—Autonomous exploration by unmanned surface vehicles (USVs) in GNSS-degraded near-shore environments is challenging. Planning must remain reliable under uneven environmental structure and localization uncertainty. The variable-resolution virtual map (VRVM) provides a scalable way to represent mapping uncertainty. However, when it is combined with a random-sampling RRT planner, the generated trajectories can be highly dispersed and may not follow the uncertainty structure encoded in the map. In this paper, we convert VRVM uncertainty into a continuous information field and use this field to guide tree expansion and trajectory proposal generation. Our proposed method guides sampling and edge expansion toward informative regions during proposal generation, and selects the best proposal using the virtual-map-based criterion to balance exploration and revisitation. Benchmarking in the VRX Gazebo simulator against VRVM with an RRT-based random-sampling planner shows that the proposed method yields more consistent exploration behavior, lower localization error, higher map coverage, and lower mapping error across the evaluated scenarios.**

***Index Terms*—Autonomous Exploration, Active SLAM, Informative Path Planning, Uncertainty-Aware Planning**

I. INTRODUCTION

Autonomous exploration by unmanned surface vehicles (USVs) is a key capability for maritime inspection, environmental monitoring, and persistent operation in near-shore waters. In such environments, however, exploration is not solely a coverage problem, because GNSS degradation, spatially uneven environmental structure, and limited onboard computation all affect how reliably the vehicle can explore while maintaining estimation quality [1]. The core issue is therefore balancing between exploration and revisitation during decision-making. A robust exploration framework should maximize spatial coverage efficiently while reducing the risk of entering information-sparse regions that degrade estimator consistency.

Recent work has addressed autonomous exploration from multiple perspectives. Some studies formulate exploration as an active-SLAM or belief-space planning problem, and evaluate candidate actions using information-theoretic objectives based on their expected effects on localization and mapping performance [2]–[15]. Other studies focus on informative planning, including next-best-view methods and sampling-or tree-based trajectory generation for information acquisition [16]–[27]. As a specialized class of information-theoretic methods, virtual map-based approaches [11]–[13], [15], [28] represent localization and mapping uncertainty using a structured uncertainty map, and evaluate candidate trajectories with respect to the uncertainty encoded in that map. However, these methods typically generate candidate trajectories through random sampling, without explicitly using the uncertainty map to guide trajectory generation. Consequently, the exploration behavior can be inconsistent and may degrade when the sampled trajectories fail to align with informative regions.

One of the latest virtual map implementations, the variable-resolution virtual map (VRVM) [28], represents mapping uncertainty using Gaussian virtual landmarks organized in an adaptive quadtree structure. While VRVM provides a strong basis for balancing exploration and revisitation, it still generates candidate trajectories using a random-sampling RRT planner. As a result, uncertainty is considered only during path evaluation, not during trajectory generation, and the resulting candidate paths can remain highly dispersed across

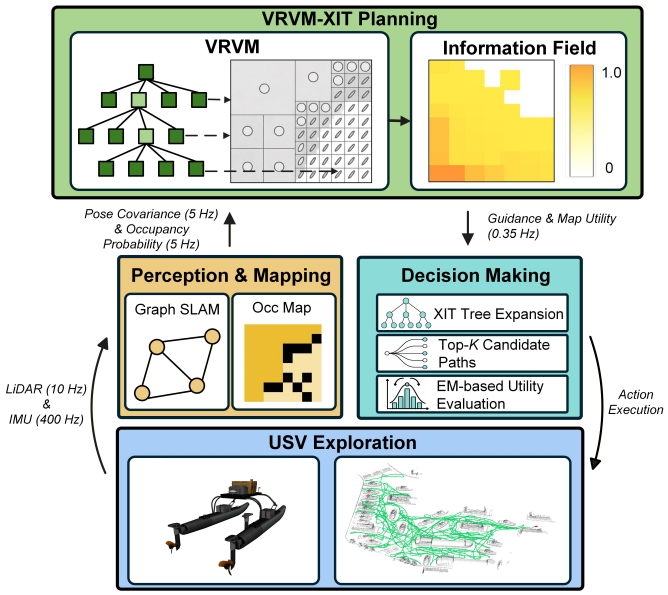

Fig. 1: Overview of the proposed framework. Graph-based LiDAR-SLAM continuously updates both the occupancy map and the variable-resolution virtual map (VRVM). The updated VRVM is then transformed into a continuous information field, which guides XIT-based tree expansion toward informative regions and supports trajectory proposal generation. The resulting proposals are evaluated through utility-based path selection, and the highest-ranked trajectory is executed by the USV.

[1]Department of Mechanical Engineering, University College London, London, UK.
[2]Department of Aeronautical and Automotive Engineering, Loughborough University, Loughborough, UK.
[3]Department of Computer Science, Dartmouth College, Hanover, NH, USA.
[4]Department of Mechanical Engineering, Stevens Institute of Technology, Hoboken, NJ, USA.
Corresponding author: Yuanchang Liu. Email: yuanchang.liu@ucl.ac.uk.

repeated runs. More broadly, existing methods emphasize either uncertainty-map representation [28] or information-aware path generation [29], while the connection between the two remains weak. A framework that jointly uses uncertainty for both trajectory generation and evaluation in autonomous exploration is still lacking.

In this paper, we bridge variable-resolution uncertainty representation and informed tree search for robust USV exploration. The uncertainty values associated with VRVM leaves are converted into a continuous information field that guides path generation by enabling waypoint and edge evaluation with respect to the surrounding information distribution. A heuristic information tree (XIT) [29] is employed to generate candidate trajectories and filter them down to a smaller set for exploration. The remaining proposals are evaluated by a utility-based selection step. The overall architecture of the proposed framework is illustrated in Fig. 1. We benchmark the proposed framework against the original VRVM implementation with RRT-based planning in an inner-port environment in the VRX Gazebo simulator. The experimental results show that the proposed algorithm achieves more consistent exploration decisions, lower translation error, and improved mapping performance.

The main contributions of this work are as follows:

- The first autonomous exploration framework that jointly incorporates localization and mapping uncertainty into both trajectory generation and trajectory evaluation.
- A principled method for converting variable-resolution virtual-map uncertainty into a continuous information field for tree-based exploration.
- An enhanced VRVM framework that achieves more consistent planning, lower translation error, and better mapping performance in GNSS-degraded marine environments.

The remainder of this paper is organized as follows. Section II presents the problem formulation. Section III describes the proposed algorithm. Section IV reports the experimental results. Section V concludes the paper.

II. Problem Formulation and Approach

We consider autonomous exploration in a previously unmapped near-shore environment. At each replanning step t, the USV maintains a SLAM-based estimate of its motion history and of the partially explored environment. Let $\mathbf{X}_{0:t} = \{\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_t\}$ denote the robot state sequence up to time t, and let $\mathbf{Z}_{0:t} = \{\mathbf{z}_0, \mathbf{z}_1, \ldots, \mathbf{z}_t\}$ denote the corresponding measurement sequence. Using a factor-graph SLAM back-end, the state-estimation problem is written as

$$\mathbf{X}^{\star}_{0:t} = \arg\max_{\mathbf{X}_{0:t}} p(\mathbf{X}_{0:t} \mid \mathbf{Z}_{0:t}), \tag{1}$$

and the latest marginal state belief is represented by the latest state $\mathbf{x}_t$ and its marginal covariance $\Sigma_{\mathbf{x}_t}$.

In parallel with SLAM, the explored workspace is represented in two complementary forms. The first is an occupancy map $\mathcal{M}_t$, which is used for collision checking and feasibility evaluation. The second is the variable-resolution virtual map. Following the VRVM formulation, we denote by $\mathbf{V}^{\star}_t$ the current deterministic estimate of the virtual map:

$$\mathbf{V}^{\star}_t \;\approx\; \arg\max_{\mathbf{V}} p(\mathbf{V} \mid \mathbf{X}^{\star}_{0:t}, \mathbf{Z}_{0:t}), \tag{2}$$

where $\mathbf{V}$ denotes the set of virtual landmarks used to represent the uncertainty associated with exploration, specifically the localization uncertainty over the explored region. We therefore write the replanning belief at time t as

$$\mathcal{B}_t = (\mathbf{x}_t,\ \Sigma_{\mathbf{x}_t},\ \mathcal{M}_t,\ \mathbf{V}^{\star}_t). \tag{3}$$

Here, $\mathcal{M}_t$ characterizes the collision-free space currently available to the USV, whereas $\mathbf{V}^{\star}_t$ describes the spatial distribution of exploration uncertainty. Different from the previous approach in [28], $\mathbf{V}^{\star}_t$ is not limited to frontier selection; instead, it serves as a structured representation of localization and mapping uncertainty for guiding path planning.

Given the current belief $\mathcal{B}_t$, the exploration problem is to select a feasible future trajectory segment

$$\mathbf{X}_{t:t+K} = \{\mathbf{x}_t, \mathbf{x}_{t+1}, \ldots, \mathbf{x}_{t+K}\},$$

where K denotes the number of future states considered in one replanning cycle. Let $\Pi(\mathcal{B}_t)$ denote the set of dynamically feasible and collision-free trajectory segments under the current belief. The replanning problem is then written as

$$\mathbf{X}^{\star}_{t:t+K} = \arg\max_{\mathbf{X}_{t:t+K} \in \Pi(\mathcal{B}_t)} U_{\mathrm{EM}}(\mathbf{X}_{t:t+K}), \tag{4}$$

with $U_{\mathrm{EM}}(\cdot)$ be the trajectory-level utility. The introduction of this trajectory-level evaluation ensures that exploration uncertainty is accounted for continuously during trajectory generation, rather than only at individual sampled locations along the path.

Since $\mathbf{V}^{\star}_t$ is a spatially discrete representation, while candidate trajectories must be generated in continuous space, we transform the latest virtual map $\mathbf{V}^{\star}_t$ into an information field:

$$\mathcal{I}_t(\mathbf{p}):\ \mathcal{X}_{\mathrm{free}}(\mathcal{M}_t) \to [0, 1], \tag{5}$$

defined over the collision-free planning space. This field must satisfy three requirements: consistency with the exploration uncertainty structure encoded by $\mathbf{V}^{\star}_t$, queryability at arbitrary planning states or along path segments, and interpretability under variable spatial resolution. Section III shows how the VRVM is converted into an information field for trajectory generation.

III. VRVM-XIT Algorithm

Starting from the latest variable-resolution virtual map $\mathbf{V}^{\star}_t$, we first construct the information field $\mathcal{I}_t(\mathbf{p})$. This field is then used to generate informed tree-based trajectory proposals. Finally, all proposals are evaluated at the path level under the current replanning belief, and the best one is selected for execution.

A. Continuous Guidance From Updated VRVM

At replanning step t, the updated virtual map $\mathbf{V}_t^\star = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{N_t}\}$ contains the active virtual landmarks maintained by the VRVM update process. Each landmark $\mathbf{v}_k$ is associated with a spatial support region and a covariance matrix $\Sigma_{\mathbf{v}_k}$ that describes the uncertainty of that region. To guide planning, we assign an uncertainty value to each position $\mathbf{p}$ in the collision-free space $\mathcal{X}_{\text{free}}(\mathcal{M}_t)$.

Let $\mathbf{v}_{k(\mathbf{p})}$ denote the virtual landmark of $\mathbf{V}_t^\star$ at position $\mathbf{p}$. We assign to $\mathbf{p}$ an uncertainty value based on the D-optimality criterion:

$$U_t(\mathbf{p}) = \log\det\big(\Sigma_{\mathbf{v}_{k(\mathbf{p})}}\big)\,. \tag{6}$$

Since the absolute scale of $\log\det(\Sigma_{\mathbf{v}_k})$ varies with the current map state, we normalize it over the active virtual landmarks:

$$U_t^{\min} = \min_{\mathbf{v}_j \in \mathbf{V}_t^\star} \log\det(\Sigma_{\mathbf{v}_j}), \qquad U_t^{\max} = \max_{\mathbf{v}_j \in \mathbf{V}_t^\star} \log\det(\Sigma_{\mathbf{v}_j}). \tag{7}$$

The continuous information field introduced in (5) is then defined as

$$\mathcal{I}_t(\mathbf{p}) = \frac{U_t(\mathbf{p}) - U_t^{\min}}{U_t^{\max} - U_t^{\min} + \varepsilon}, \qquad \mathcal{I}_t(\mathbf{p}) \in [0,1], \tag{8}$$

where $\varepsilon > 0$ is a small constant used only to avoid numerical degeneracy when the current map uncertainty becomes nearly uniform.

For a path segment or tree edge e, we also define an edge-level field response:

$$\mathcal{I}_t(e) \approx \max_{m=0,\ldots,M_e} \mathcal{I}_t\Big(\mathbf{p}_e^{(m)}\Big)\,, \tag{9}$$

where $\{\mathbf{p}_e^{(m)}\}_{m=0}^{M_e}$ are sampled points on e. This max-pooled form avoids diluting narrow but informative regions when a segment traverses mostly low-information space.

B. Exploration-Exploitation Informed Tree Expansion

We then use $\mathcal{I}_t(\mathbf{p})$ to guide the proposal generation. Let $\mathcal{T}_t$ denote the tree rooted at the current state $\mathbf{x}_t$. At each iteration, a random sample is first drawn from free space. To avoid both purely unguided expansion and overly local greedy growth, the sampling process uses a mixed strategy: one portion of samples is retained to preserve exploratory expansion toward previously unseen free space, while the remaining samples are accepted with probability biased by $\mathcal{I}_t(\mathbf{p})$. In this way, the tree retains global exploratory capability while favoring regions that are currently assessed as informative by the updated VRVM.

After a sample is accepted, the planner searches for candidate parent nodes within a local neighborhood and evaluates feasible extensions toward the sample. Each candidate extension must satisfy collision-free motion in $\mathcal{M}_t$ and the constraints of the current kinematic model. For a feasible edge e connecting two adjacent nodes, we define the edge cost

$$C_t(e) = \alpha_d\,\ell(e) + \beta_i\,\ell(e)\big(1 - \mathcal{I}_t(e)\big), \tag{10}$$

where $\ell(e)$ is the Euclidean distance between the endpoints of edge, and $\alpha_d, \beta_i > 0$ are distance and information weights, respectively. The parent that yields the smallest accumulated path cost is selected. A local rewiring step is then applied to reduce the cost of neighboring branches when a lower-cost connection becomes available.

By repeatedly expanding the tree until a leaf node is reached, the planner constructs a set of feasible candidate trajectories $\mathcal{C}_t = \{\pi_1, \pi_2, \ldots, \pi_{M_t}\}$, where each π_i is a collision-free path from the current state to the leaf. Since the total path cost scales with path length, using it alone can bias the ranking toward shorter proposals regardless of their informative value. We therefore rank candidates using a path-level cost density:

$$J_t(\pi) = \frac{C_t(\pi)}{L(\pi) + \varepsilon}, \qquad C_t(\pi) = \sum_{e \in \pi} C_t(e), \tag{11}$$

where $L(\pi)$ denotes the path length. The candidates are sorted in ascending order of $J_t(\pi)$, and only the top-K subset $\mathcal{C}_t^K \subset \mathcal{C}_t$, with $|\mathcal{C}t^K| = K$, is retained for final evaluation. This stage only filters proposals; it does not determine the final executed trajectory.

C. Path Utility Evaluation and Replanning

The best trajectory is selected from $\mathcal{C}_t^K$ according to Eq. (4), where the utility is defined as

$$U_{\text{EM}}(\pi) = \beta\,U_{\text{pose}}(\pi) + \gamma\,U_{\text{map}}(\pi) - \alpha\,L(\pi), \tag{12}$$

$\alpha, \beta, \gamma > 0$ are weighting parameters. Here $L(\pi)$ is the total path length, and collision-infeasible proposals are discarded before scoring. The pose term, $U_{\text{pose}}(\pi)$, evaluates the predicted pose-estimation quality at the end of the proposal. Starting from the current marginal covariance $\Sigma_{\mathbf{x}_t}$, we propagate pose uncertainty along π with a local motion-covariance model and obtain the terminal covariance $\Sigma_{\mathbf{x}}^{\text{end}}(\pi)$. We then define

$$U_{\text{pose}}(\pi) = -\log\det\big(\Sigma_{\mathbf{x}}^{\text{end}}(\pi)\big)\,. \tag{13}$$

The map term, $U_{\text{map}}(\pi)$, evaluates the predicted contribution of the proposal to the reduction of virtual-map uncertainty. This quantity is approximated by sampling multiple locations along the path at a fixed arc-length interval and accumulating the predicted reduction in uncertainty within the current sensor range:

$$U_{\text{map}}(\pi) \approx \sum_{\mathbf{p} \in \mathcal{S}_{\Delta s}(\pi)} \Delta_{\text{VM}}(\mathbf{p}), \tag{14}$$

where $\mathcal{S}_{\Delta s}(\pi)$ denotes the set of path samples generated at arc-length spacing Δs, and $\Delta_{\text{VM}}(\mathbf{p})$ is the predicted uncertainty reduction queried from the updated virtual map. This formulation keeps the final path evaluation consistent with the original VRVM representation while allowing utility evaluation over a continuous trajectory.

After selection, the raw path is post-processed by shortcutting and smoothing before execution. The planner then follows the active trajectory until the current goal is reached or a replanning timeout is triggered, at which point the belief $\mathcal{B}_t$ is updated from the latest SLAM and map state and the same procedure is repeated. By converting the virtual map into a continuous guidance field, our proposed method uses the localization and mapping uncertainty criteria to generate and evaluate proposals at the trajectory level.

IV. EXPERIMENTS AND RESULTS

A. Experimental Setup

We conducted autonomous exploration experiments in a simulated docker-basin inner-port environment of approximately $210\,\mathrm{m}\times500\,\mathrm{m}$, populated with densely distributed ship-shaped obstacles of varying scales and local feature densities. All experiments were performed in the same Gazebo [30] framework, using a common LiDAR–IMU sensing configuration, the same LIO-SAM back-end for pose and covariance estimation [31], and the same projected occupancy map for collision checking. The virtual map was updated online within a fixed workspace of $[-300, 300]\,\mathrm{m}\times[-300, 300]\,\mathrm{m}$, with base resolution $1.0\,\mathrm{m}$, prior standard deviation $\sigma_0 = 2.0$, sensor range $30.0\,\mathrm{m}$, and covariance-intersection fusion enabled using the log-determinant objective. Unless otherwise stated, the VRVM parameters were fixed to $\tau_{0,\mathrm{det}} = 1.0$, $\tau_{0,\mathrm{prob}} = 0.15$, and $\texttt{max\_depth} = 7$.

The proposed method, VRVM-XIT, was compared against the previous VRVM planner with RRT-based proposal generation, hereafter denoted VRVM-RRT. Both methods use the same perception, occupancy, and virtual-map back-end, but differ in how trajectory proposals are generated. To ensure a fair comparison, both planners were configured with matched proposal budgets and comparable utility scaling. In particular, both planners used 50 candidate paths, 3000 tree-expansion iterations, a step length of $5.0\,\mathrm{m}$, minimum and maximum path lengths of $3.0\,\mathrm{m}$ and $80.0\,\mathrm{m}$, respectively, and the same sensing range. The path-length penalty, localization term, and mapping term were also tuned to remain on comparable numerical scales in the two planners, with matched utility weights $(\alpha, \beta, \gamma) = (0.1,\, 0.1,\, 10^{-4})$. This alignment is important because the comparison is intended to isolate the effect of uncertainty-guided proposal generation and action selection, rather than to favor one method through unequal utility weighting.

The experiments are organized in three parts. First, to study the coupling between VRVM representation and planning behavior, we vary the VRVM refinement depth over $\texttt{max\_depth} \in \{5, 6, 7, 8\}$ and visualize the resulting exploration process, information field, and translation error over the first 20 exploration steps. Second, to evaluate local planning consistency, we select three local docker-basin scenes with different free-space and feature configurations and restart VRVM-XIT and VRVM-RRT 30 times from the same initial condition in each scene. From these repeated trials, we report Goal Endpoint Dispersion (GED). Finally, we perform one representative extended exploration run for each planner in a $210\,\mathrm{m} \times 250\,\mathrm{m}$ basin sector and compare Union-Normalized Mapping Coverage (UNC), Ground-Truth-Consistent Mapping Error (GME), and translation error over time.

B. Planning Consistency Across Local Scenes

We next evaluate the planning consistency of VRVM-XIT and VRVM-RRT over three local scenes extracted from the docker-basin environment. These scenes share the same global map and planner configuration, but differ in the amount and directional distribution of free space visible from the start pose. In Group 1, the robot initially observes the largest collision-free region, so multiple exploration directions are immediately admissible. In Group 2, the initial free space is smaller, but local features remain visible in all directions. In Group 3, only approximately one half of the initial view contains feasible free space, making this the most directionally constrained case. For each scene, both planners were restarted independently 30 times from the same initial pose, and the first valid path generated in each trial was recorded. We then quantify planning consistency using the dispersion of path endpoints with respect to the centroid of the endpoint set.

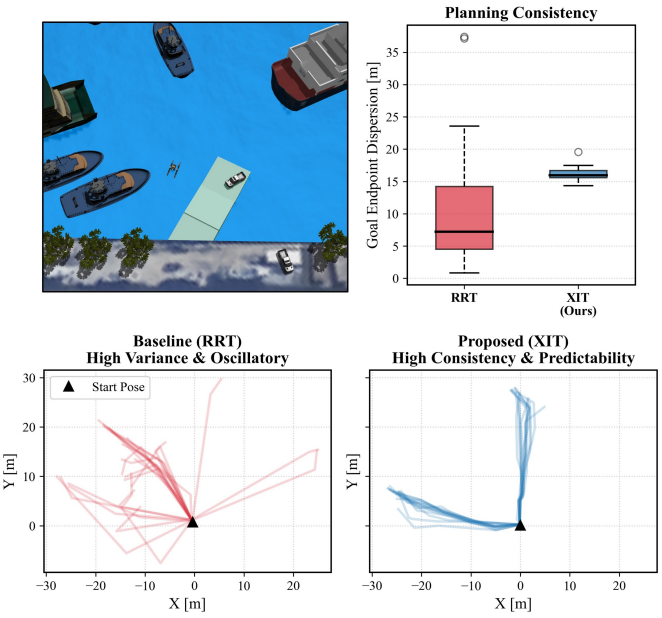

Fig. 2: Planning consistency in Group 1, where the initial collision-free region is the largest and multiple exploration directions are simultaneously admissible. The top row shows the local scene and the Goal Endpoint Dispersion (GED) comparison over 30 independent restarts for VRVM-RRT and VRVM-XIT. The bottom row overlays the first valid planned path from each restart. In this multi-modal case, VRVM-XIT forms a small number of structured decision modes rather than a single endpoint cluster, whereas VRVM-RRT produces more scattered and oscillatory proposals.

The corresponding results for Groups 1, 2, and 3 are shown in Fig. 2, Fig. 3, and Fig. 4, respectively. In all three scenes, the baseline VRVM-RRT produces visibly scattered path families, with substantial run-to-run variation in both heading and path geometry. By contrast, VRVM-XIT generates more structured and repeatable path clusters. This difference is especially clear in Groups 2 and 3, where the endpoint-dispersion boxplots show a clear reduction relative to RRT, and the overlaid paths indicate that XIT repeatedly selects the same admissible exploration directions. Group 1 is less straightforward. Because this scene contains the largest initial free-space region, several early actions are simultaneously feasible, and XIT tends to form a small number of consistent decision modes rather than a single endpoint cluster. As a result, the centroid-based dispersion metric is less favorable to XIT in this

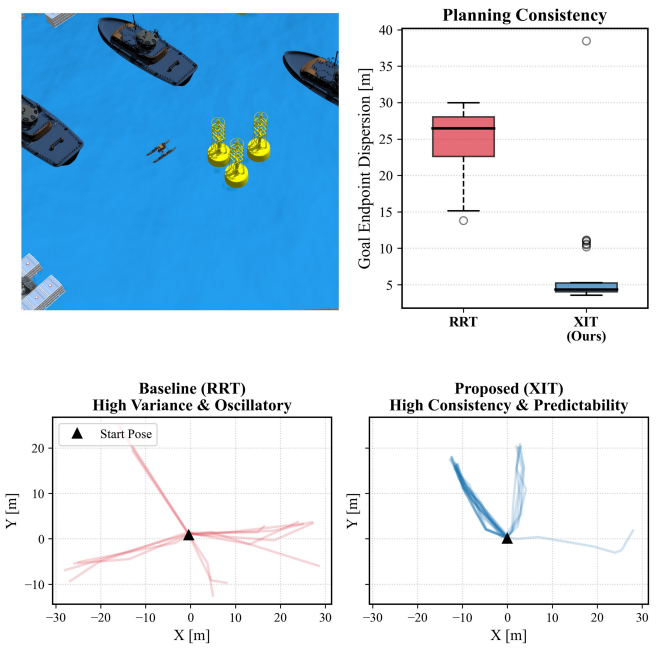

Fig. 3: Planning consistency in Group 2, where the initial free space is smaller but local structural features remain visible in all directions. The top row shows the local scene and the GED comparison over 30 independent restarts for VRVM-RRT and VRVM-XIT. The bottom row overlays the first valid planned path from each restart. Compared with the RRT-based baseline, VRVM-XIT produces more repeatable path families and lower endpoint dispersion, indicating more stable local action selection under the same initial condition.

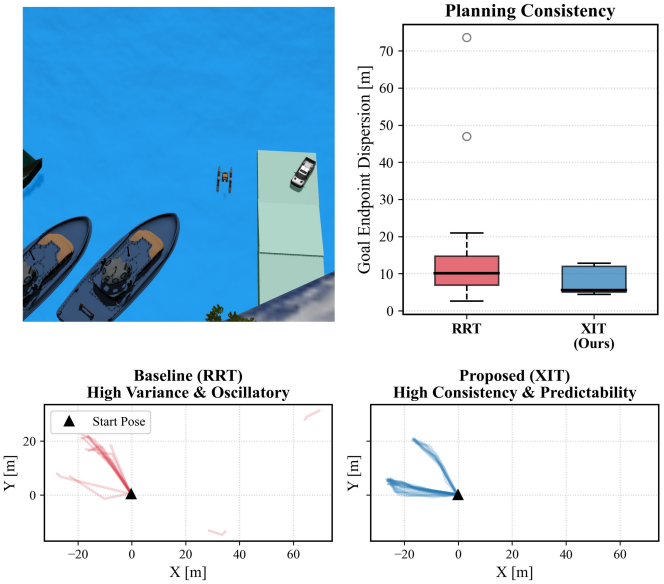

Fig. 4: Planning consistency in Group 3, the most directionally constrained local scene, where only approximately one half of the initial view contains feasible free space. The top row shows the local scene and the GED comparison over 30 independent restarts for VRVM-RRT and VRVM-XIT. The bottom row overlays the first valid planned path from each restart. In this strongly constrained setting, VRVM-XIT repeatedly selects similar admissible directions, while VRVM-RRT remains substantially more dispersed across repeated runs.

case, even though the overlaid trajectories remain noticeably more organized than the highly oscillatory RRT proposals. Overall, these results indicate that the proposed uncertainty-guided planner produces more predictable local exploration actions than the RRT-based baseline. This advantage becomes more evident when, under the same localization condition and within the same locally reachable free space, multiple feasible exploration directions are available. In such cases, VRVM-XIT more consistently drives the planned path toward locally preferred informative directions, whereas VRVM-RRT produces more dispersed trajectories due to random proposal generation.

C. Extended Exploration Behavior, Mapping Quality, and Translation Error

We finally compare VRVM-XIT and VRVM-RRT in an extended exploration task over a representative $210\,\mathrm{m} \times 250\,\mathrm{m}$ basin sector. The two planners share the same initial pose, the same GT evaluation region, the same VRVM configuration with `max_depth` $= 7$, and the same matched utility scaling for the mapping, localization, and path-length terms. For this representative comparison, each planner was run once under the same exploration framework until the explored map approached the target GT coverage level. We then evaluate three task-level quantities: union-normalized map coverage over the aligned ROI, GT-consistent map RMSE, and translation error over time. The coverage and mapping-error curves are computed from the aligned projected occupancy maps using the union ROI and GT distance-field procedures implemented in our evaluation scripts.

The quantitative results in Fig. 6 and Fig. 7 show a clear advantage of VRVM-XIT in this representative comparison. In terms of coverage growth, both methods expand the explored region at a comparable rate at the beginning of the run, but the proposed planner maintains useful progress for longer and continues to increase coverage after VRVM-RRT has largely entered a plateau. Under the same target-coverage setting, this yields a more sustained coverage-growth profile for VRVM-XIT, which reaches approximately 0.95 coverage, whereas VRVM-RRT levels off near 0.90. At the same time, the higher coverage is not obtained at the expense of map quality. The map RMSE of VRVM-XIT remains lower throughout most of the run and stabilizes at roughly 1.3–$1.4\,\mathrm{m}$, while VRVM-RRT rises to about $2.1\,\mathrm{m}$. This representative result indicates a better coverage–accuracy trade-off for VRVM-XIT in the reported long-horizon exploration setting. The translation-error comparison shows the same trend: the mean translation error of VRVM-XIT is approximately $0.73\,\mathrm{m}$, compared with about $1.19\,\mathrm{m}$ for VRVM-RRT.

Figure 5 qualitatively illustrates the corresponding exploration behavior. The VRVM-RRT trajectory exhibits dense recrossing and more redundant traversals before reaching the target coverage level, whereas the VRVM-XIT trajectory is more spatially organized and expands through the basin

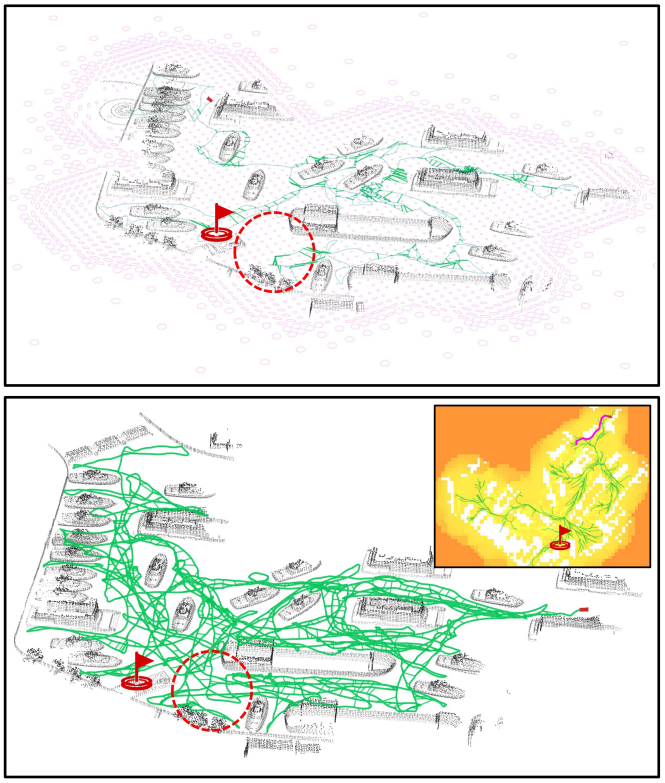

Fig. 5: Representative long-horizon exploration behavior in the $210\,\mathrm{m} \times 250\,\mathrm{m}$ basin sector. The top panel shows the exploration trajectory of VRVM-RRT, which exhibits denser recrossing and more redundant traversals before approaching the target coverage level. The bottom panel shows the corresponding result of VRVM-XIT, whose trajectory is more spatially organized and expands through the basin sector in a more directed manner; the inset further shows that the late-stage information field still provides coherent guidance for local proposal generation.

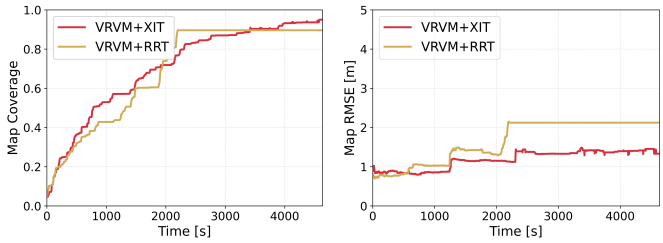

Fig. 6: Representative long-horizon comparison of mapping performance in the basin-sector experiment. Left: Union-Normalized Mapping Coverage (UNC) over time. Right: Ground-Truth-Consistent Mapping Error (GME) over time. Both methods achieve comparable early coverage growth, but VRVM-XIT maintains useful progress for longer and reaches a higher late-stage coverage level, while also preserving lower map error throughout most of the run.

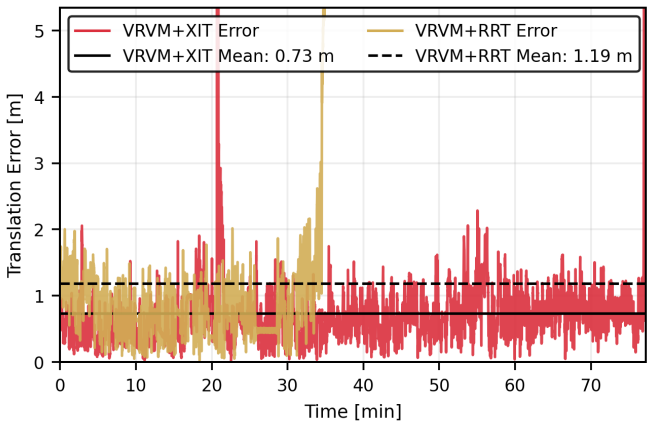

Fig. 7: Representative long-horizon comparison of translation error over time in the basin-sector experiment. Under the same VRVM back-end, sensing configuration, and matched utility scaling, VRVM-XIT yields lower translation error than VRVM-RRT during most of the reported run.

sector in a more directed manner. The red-circled region in Fig. 5 illustrates this difference more clearly. Even at this late exploration stage, the updated VRVM still highlights nearby regions with higher expected value for localization support. Guided by this information, VRVM-XIT selects a revisitation direction that is more consistent with localization correction and possible loop closure, whereas VRVM-RRT, due to its random proposal generation, may instead return to previously visited areas whose revisitation is less informative for the current estimation state. This observation is consistent with the mechanism established in Sections III–IV-B: by converting the current variable-resolution uncertainty map into continuous planning guidance, VRVM-XIT yields more stable local decisions in the presented experiments, and these local advantages accumulate into better long-horizon exploration behavior in the reported basin-sector comparison.

V. CONCLUSIONS

This paper presented VRVM-XIT, an uncertainty-guided exploration framework for unmanned surface vehicles operating in GNSS-degraded near-shore environments. The method uses a variable-resolution virtual map not only as a scalable representation of mapping uncertainty, but also as the basis for constructing a continuous guidance field for planning. This improves the connection between uncertainty representation and action generation. On this basis, VRVM-XIT combines informed tree expansion with Top-K proposal filtering and utility-based trajectory selection. Experimental comparisons against a VRVM-RRT baseline showed more consistent local exploration decisions, lower translation error, and better mapping performance in the evaluated simulation scenarios.

ACKNOWLEDGMENT

The authors acknowledge the support of the Engineering and Physical Sciences Research Council (EPSRC) under Grants EP/Y000862/1. Y. Liu is also supported by The Royal Society Kan Tong Po Fellowship (KTP/R1/251117).

REFERENCES

[1] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023.

[2] A. Kim and R. M. Eustice, "Active visual slam for robotic area coverage: Theory and experiment," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 457–475, 2015. [Online]. Available: https://journals.sagepub.com/doi/abs/10.1177/0278364914547893

[3] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments," *International Journal of Robotics Research*, vol. 34, no. 7, pp. 849–881, 2015. [Online]. Available: https://journals.sagepub.com/doi/abs/10.1177/0278364914561102

[4] H. Carrillo, P. M. Dames, V. Kumar, and J. A. Castellanos, "Autonomous robotic exploration using occupancy grid maps and graph SLAM based on shannon and rényi entropy," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 487–494.

[5] H. Carrillo, Y. Latif, M. L. Rodríguez-Arévalo, J. Neira, and J. A. Castellanos, "On the monotonicity of optimality criteria during exploration in active SLAM," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1476–1483.

[6] P. Morere, R. Marchant, and F. Ramos, "Sequential bayesian optimization as a pomdp for environment monitoring with uavs," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017.

[7] M. L. Rodríguez-Arévalo, J. Neira, and J. A. Castellanos, "On the importance of uncertainty representation in active slam," *IEEE Transactions on Robotics*, vol. 34, no. 3, pp. 829–834, 2018.

[8] Y. Chen, S. Huang, R. Fitch, L. Zhao, H. Yu, and D. Yang, "On-line 3d active pose-graph SLAM based on key poses using graph topology and sub-maps," in *2019 IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 169–175.

[9] J. A. Placed and J. A. Castellanos, "Fast autonomous robotic exploration using the underlying graph structure," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 6672–6679.

[10] ——, "A general relationship between optimality criteria and connectivity indices for active graph-slam," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 816–823, 2023.

[11] J. Wang and B. Englot, "Autonomous exploration with expectation-maximization," in *Robotics Research: The 18th International Symposium ISRR*. Springer, 2020, pp. 759–774.

[12] J. Wang, T. Shan, and B. Englot, "Virtual maps for autonomous exploration with pose slam," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4899–4906.

[13] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan, and B. Englot, "Virtual maps for autonomous exploration of cluttered underwater environments," *IEEE Journal of Oceanic Engineering*, vol. 47, no. 4, pp. 916–935, 2022.

[14] Y. Huang, X. Lin, and B. J. Englot, "Multi-robot autonomous exploration and mapping under localization uncertainty with expectation-maximization," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 7236–7242.

[15] I. Collado-Gonzalez, J. McConnell, J. Wang, P. Szenher, and B. Englot, "Real-time planning under uncertainty for auvs using virtual maps," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 8334–8340.

[16] S. Bai, J. Wang, F. Chen, and B. Englot, "Information-theoretic exploration with bayesian optimization," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 1816–1822.

[17] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon "next-best-view" planner for 3d exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1462–1468.

[18] M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake, "Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring," *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 658–685, 2019.

[19] N. Palomeras, N. Hurtós, E. Vidal, and M. Carreras, "Autonomous exploration of complex underwater environments using a probabilistic next-best-view planner," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1619–1625, 2019.

[20] M. Selin, M. Tiger, D. Duberg, F. Heintz, and P. Jensfelt, "Efficient autonomous exploration planning of large-scale 3-d environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1699–1706, 2019.

[21] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.

[22] M. Popović, T. Vidal-Calleja, J. J. Chung, J. Nieto, and R. Siegwart, "Informative path planning for active field mapping under localization uncertainty," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 751–10 757.

[23] C. Rhodes, C. Liu, and W.-H. Chen, "Informative path planning for gas distribution mapping in cluttered environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 6726–6732.

[24] H. Zhu, J. J. Chung, N. R. Lawrance, R. Siegwart, and J. Alonso-Mora, "Online informative path planning for active information gathering of a 3d surface," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1488–1494.

[25] J. Ott, E. Balaban, and M. J. Kochenderfer, "Sequential bayesian optimization for adaptive informative path planning with multimodal sensing," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7894–7901.

[26] J. Rückin, F. Magistri, C. Stachniss, and M. Popović, "An informative path planning framework for active learning in uav-based semantic mapping," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4279–4296, 2023.

[27] W. Chen, R. Khardon, and L. Liu, "Adaptive robotic information gathering via non-stationary gaussian processes," *The International Journal of Robotics Research*, vol. 43, no. 4, pp. 405–436, 2024.

[28] Y. Li, Y. Huang, W. GaoZhang, A. Quattrini Li, B. Englot, and Y. Liu, "Variable-resolution virtual maps for autonomous exploration with unmanned surface vehicles (usvs)," *arXiv preprint arXiv:2603.22667*, 2026. [Online]. Available: https://arxiv.org/abs/2603.22667

[29] M. Fazliu, M. Coombes, S. Wang, and C. Liu, "Xit: Exploration and exploitation informed trees for active gas distribution mapping in unknown environments," *arXiv preprint arXiv:2602.13739*, 2026. [Online]. Available: https://arxiv.org/abs/2602.13739

[30] B. Saldarriaga-Mesa, J. Montesdeoca, D. Báez, F. Roberti, and J. M. Toibero, "Open-access simulation platform and motion control design for a surface robotic vehicle in the VRX environment," *Robotics*, vol. 14, no. 10, 2025.

[31] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5135–5142.