Variable-Resolution Virtual Maps for Autonomous Exploration with Unmanned Surface Vehicles

Ye Li[1], Yewei Huang[2], Wenlong GaoZhang[1],
Alberto Quattrini Li[2], Brendan Englot[3], and Yuanchang Liu[1]
[1]University College London [2]Dartmouth College [3]Stevens Institute of Technology

Introduction

Autonomous exploration by unmanned surface vehicles (USVs) in near-shore waters requires reliable localisation and consistent mapping over extended areas, but this is challenged by GNSS degradation, environment-induced localisation uncertainty, and limited on-board computation. We propose a Variable-Resolution Virtual Map (VRVM), a computationally efficient method for representing map uncertainty using bivariate Gaussian virtual landmarks placed in the cells of an adaptive quadtree. We evaluate VRVM in the VRX Gazebo simulator, using a realistic marina environment. We summarise our **main contributions** as follows:

- To the best of our knowledge, VRVM is the first virtual map implementation to achieve real-time performance that is compatible with **embedded systems**.
- An area-weighted map valuation strategy balances **exploration and exploitation** in structurally unbalanced nearshore coastal environments.

VRVM Data Structure

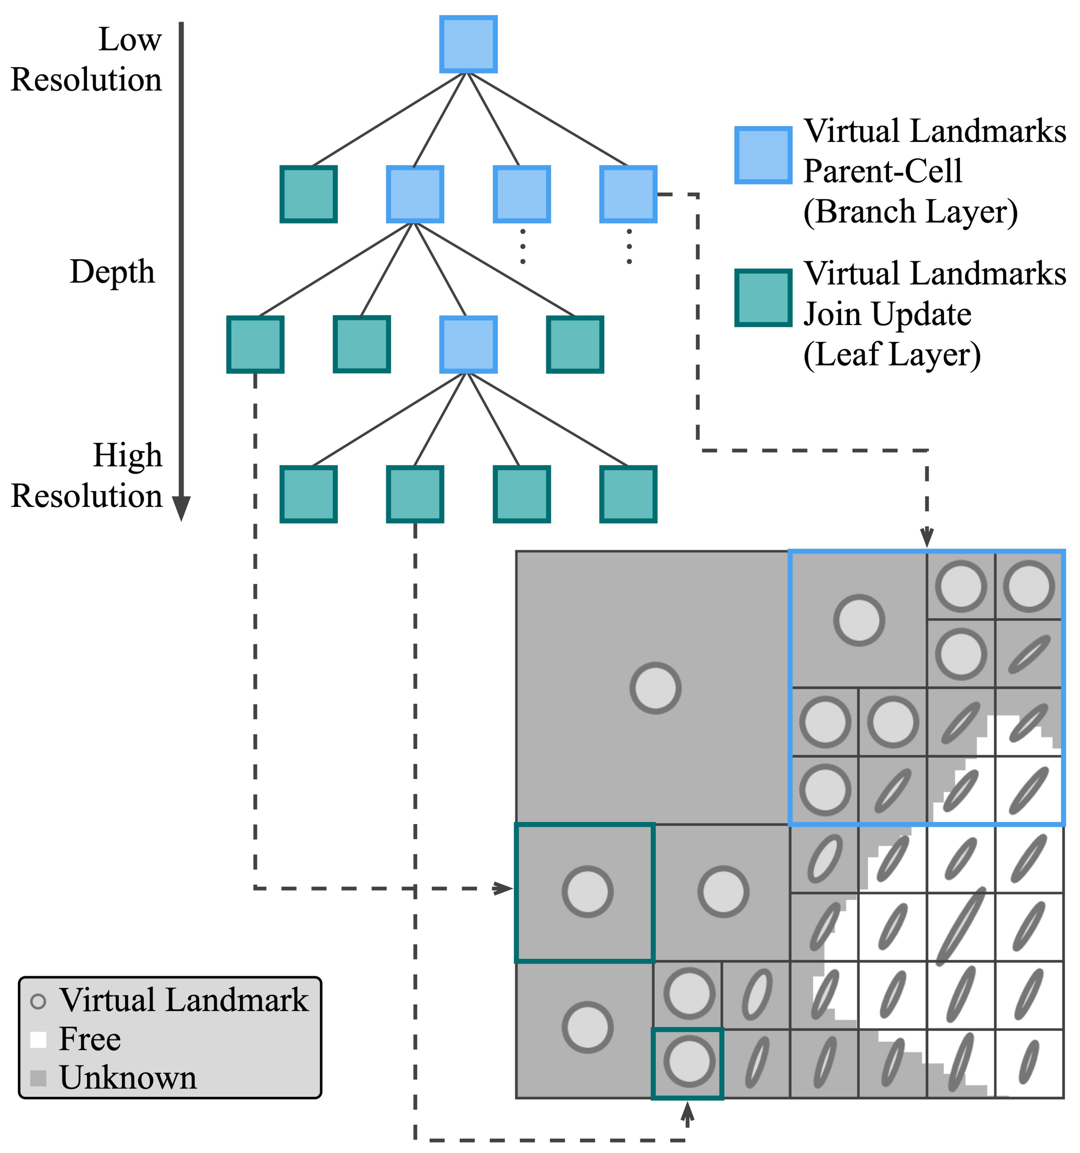

Figure 1: Visualization of a portion of the variable-resolution virtual map hierarchy over a quadtree. In this quadtree $\mathscr{V}$, each leaf $\mathbf{v}_q$ represents a virtual landmark and carries an area indicator specifying the spatial coverage of the leaf. During exploration, the quadtree is refined adaptively to ensure that the virtual map provides the resolution needed for planning.

Computational Efficiency

Performance of VRVM in the port-area scene on desktop host and **Raspberry Pi** platforms. The top row shows Resident Set Size (RSS) memory usage relative to a 450 MB budget, while the bottom row shows planning step-gap intervals against a 10 s failure threshold.

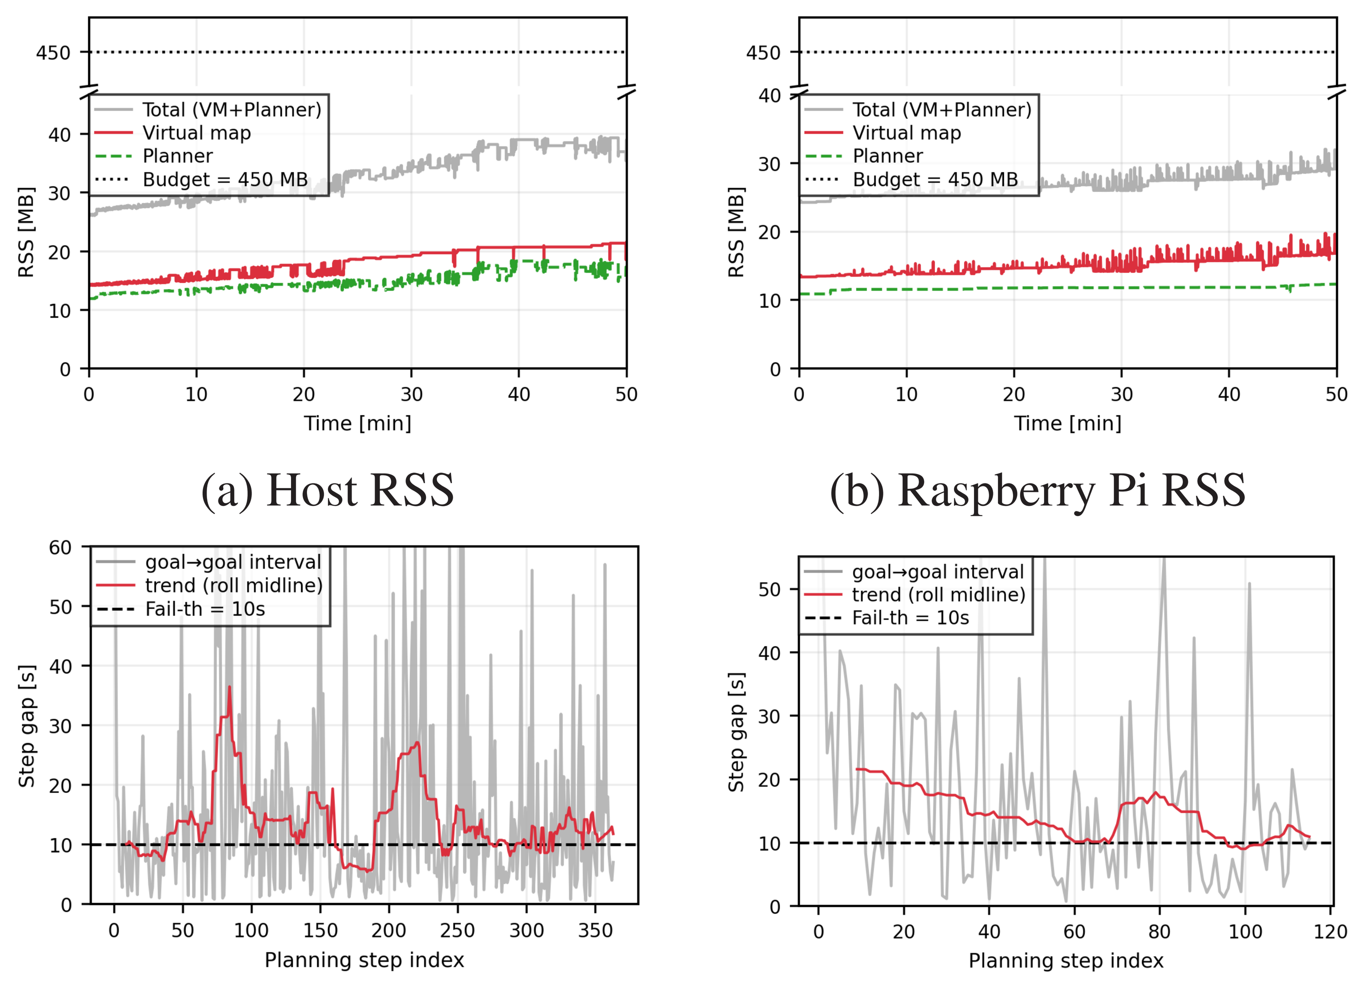

(a) Host RSS (b) Raspberry Pi RSS
(c) Host step gap (d) Raspberry Pi step gap

Exploration using Variable-Resolution Virtual Map (VRVM)

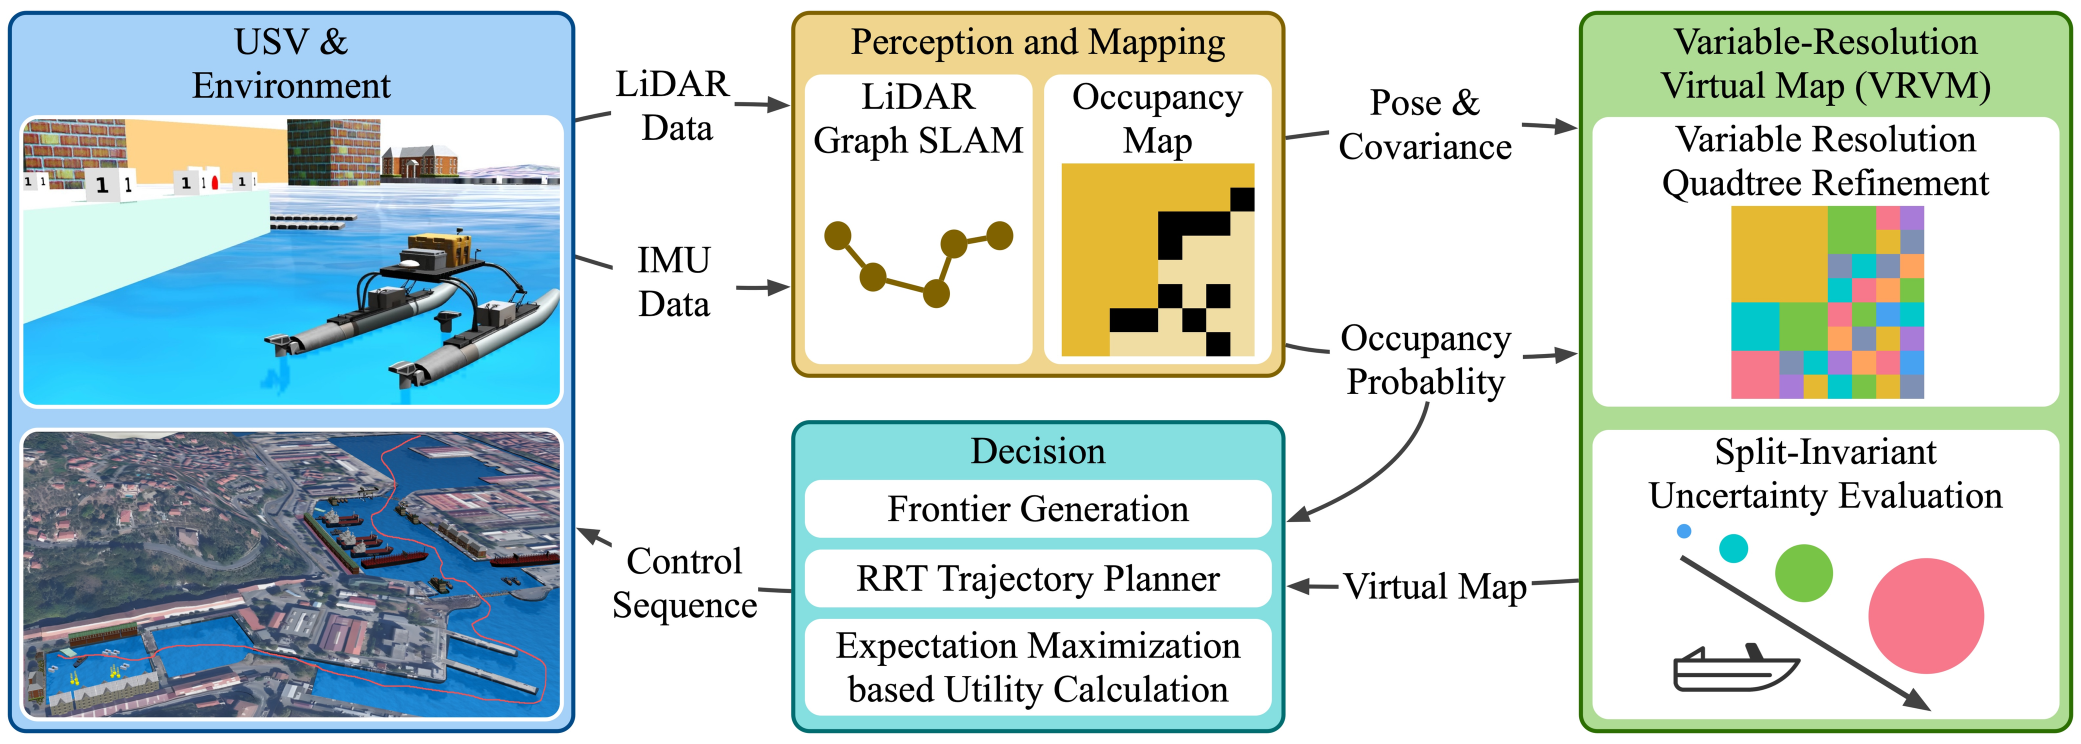

Figure 2: Pipeline overview. LiDAR and IMU data are processed by a graph-SLAM backend to estimate the USV trajectory, update the occupancy map, and provide pose covariance. These outputs update the **Variable-Resolution Virtual Map (VRVM)**, which uses adaptive quadtree refinement and area-weighted uncertainty evaluation to represent map uncertainty efficiently. Frontier candidates and feasible trajectories are then evaluated by an **Expectation-Maximisation (EM)**-style planner using pose uncertainty, VRVM-based map uncertainty reduction, and path length cost. The selected trajectory is executed in a receding-horizon loop, enabling **uncertainty-aware exploration** with improved computational scalability.

Split-Invariant Uncertainty Valuation

Directly summing log-determinants over leaves is split-sensitive: subdividing a coarse region can change the score even when local uncertainty is unchanged. VRVM uses area-normalized weights:

$$w_{\mathrm{area}}(\mathbf{v}_q,\mathscr{V}) = \frac{A_q}{\sum_{\mathbf{v}_r\in\mathscr{V}} A_r}, \qquad J_{\mathrm{area}}(\mathscr{V}) = \sum_{\mathbf{v}_q\in\mathscr{V}} w_{\mathrm{area}}(\mathbf{v}_q,\mathscr{V})\,\phi(\mathbf{v}_q).$$

For a candidate control sequence π, the map term combines the predicted end-of-horizon map uncertainty and the discounted visible-set gain:

$$U_{\mathrm{map}}(\pi) = J_{\mathrm{area}}(\mathscr{V}_{t+h}) + \Delta J_{\mathrm{gain}}(\pi), \qquad \Delta J_{\mathrm{gain}}(\pi) = \sum_{k=0}^{t+h}\left[\gamma_k J_{\mathrm{area}}(\mathscr{V}_{t+h}(\mathbf{x}_k)) - J_{\mathrm{area}}(\mathscr{V}_t(\mathbf{x}_k))\right].$$

This valuation keeps coarse, feature-sparse regions deliberately uncertain while preserving sensitivity to information-dense structure.

Representative Results

We run all experiments in the VRX simulation environment using the standard WAM-V USV platform equipped with a 16-beam 3D LiDAR and an IMU. To capture the structural diversity of real near-coastal waters, we use a large, composite water-area setting. For state estimation, we use LIO-SAM with GNSS inputs disabled; its 6-DoF estimates are projected onto the horizontal plane to form the planner's $SE(2)$ state. In parallel, we integrate the registered point clouds into a fixed-resolution 3D OctoMap and then project it to a 2D occupancy grid.

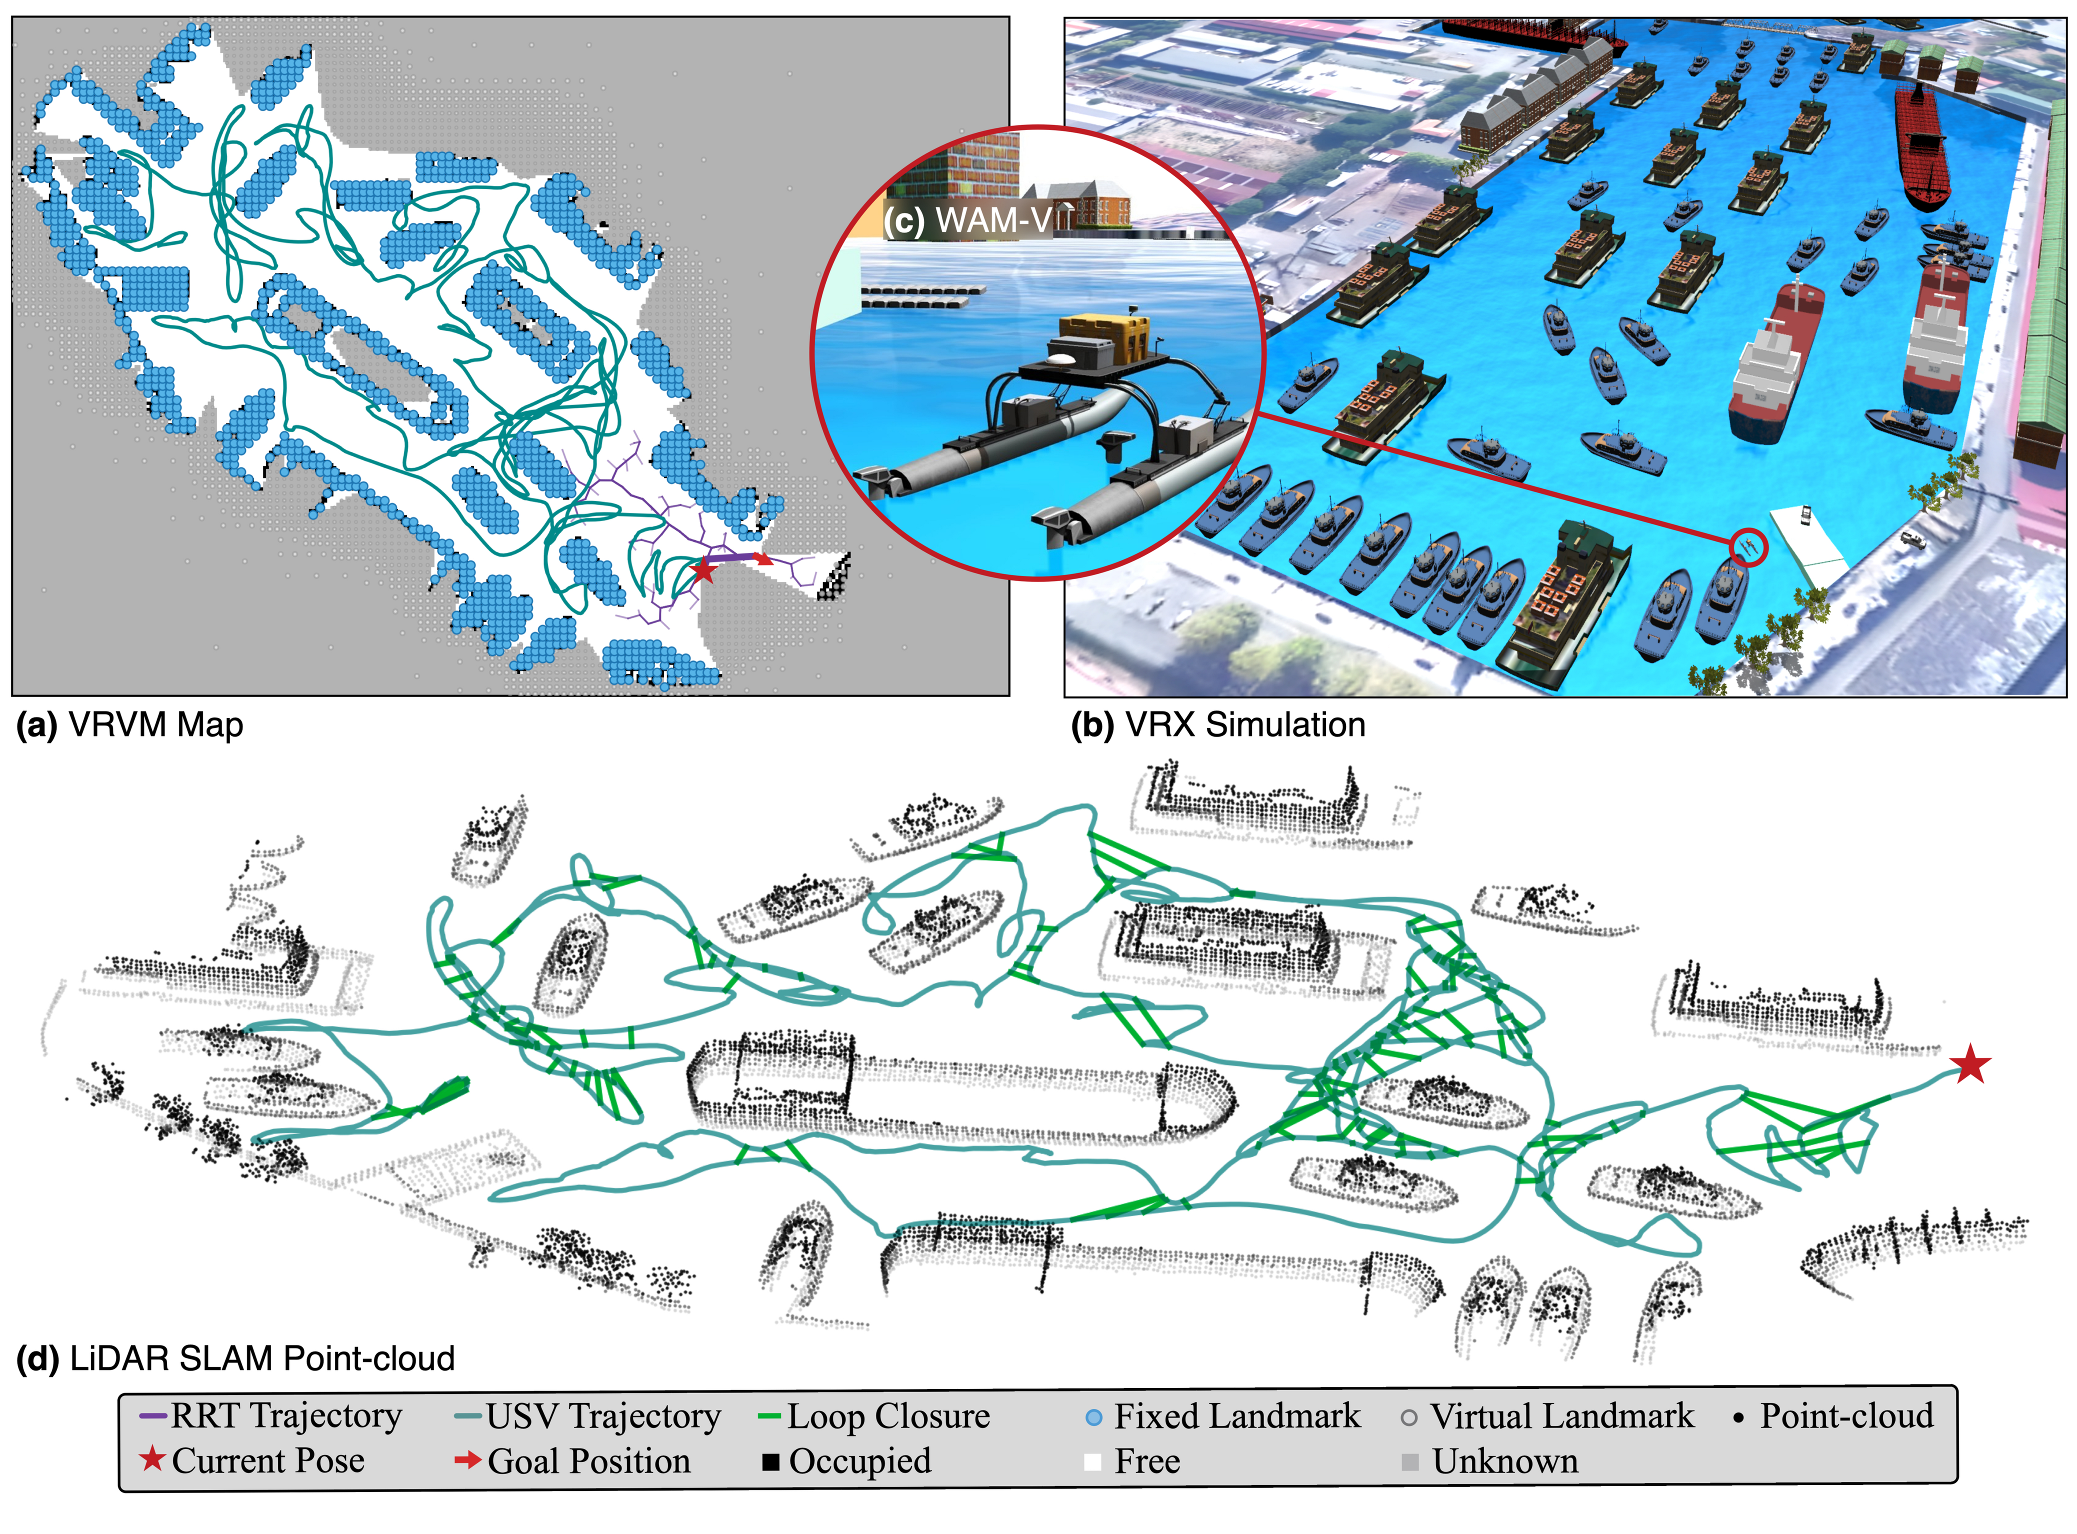

Figure 3: Illustration of the WAM-V (c) in the 210 m × 500 m Harbour Basin scene (b). The VRVM map (a) and corresponding LiDAR SLAM point cloud (d) are shown. The USV is depicted in red, with the trajectory history in green. White cells indicate observed space, and gray cells denote unobserved space. Ellipses represent the covariance associated with each virtual landmark.