Mission Planning Framework for Autonomous Marine Vehicles in Seabed Surveys

S.H. Ahmed
UiT The Arctic University of Norway, Norway

L.P. Perera
UiT The Arctic University of Norway, Tromsø, Norway
SINTEF Digital, SINTEF AS, Oslo, Norway

M. Adnan
UiT The Arctic University of Norway, Norway

ABSTRACT: Seabed surveying with surface and underwater vehicles increasingly relies on cooperative unmanned platforms that must remain safe, COLREG-compliant, and transparent to remote operators. This paper reviews mission-planning for surface and underwater vehicles, and presents an auditable layered architecture for an onshore-supervised surface vessel that generates a coverage plan from an operator-defined polygon and produces structured reports for autonomous underwater vehicle inspection. A behavior-tree executive that sequences mission phases and handles events are presented in this study. A coverage planner generates a baseline survey backbone and applies bounded local edits, supported by a reinforcement-learning advisor that recommends when and where edits are beneficial without direct actuation. A local tracking controller executes the active segment, while a risk assessment and collision-avoidance layer enforce safety and logs overrides. The proposed framework is grounded in the ongoing surface vessel and remote operation center (ROC) development, enabling incremental validation through repeatable logs and status reporting.

1 INTRODUCTION

Autonomous maritime surveying is increasingly framed as a cooperative operation in which an unmanned surface vessel (USV) provides wide-area sensing and coordination while underwater vehicles (AUVs) conduct closer inspection and data fusion. In seabed survey missions, the USV can rapidly map large areas using sensors such as multibeam echosounders (MBES), then generate structured target reports that support follow-up inspection by one or more AUVs. These frameworks are especially relevant in underwater object detection and persistent underwater infrastructure monitoring contexts, where sensing dropouts, intermittent links, and dynamic encounter situations are common, and where robust recovery behaviors can be mission critical (Thompson and Guihen, 2019, Ferri and Djapic, 2013) due to rough weather conditions.

Despite rapid progress, designing a field-ready mission planner remains challenging because maritime autonomy is rarely a single algorithmic choice. In practice, a remote operation center (ROC) is needed to supervise the USV, maintain situational awareness, and provide a human-in-the-loop safety backstop when unexpected events occur (Adnan et al., 2024b). The mission therefore must be expressed in a form that supports phase sequencing, failure handling, and clear supervision from the ROC, so that the operator can monitor progress, understand system state, and intervene through well-defined emergency procedures when required (Adnan and Perera, 2025). Functionally, the planner must transform a survey polygon into an efficient coverage plan (Zhao et al., 2024), execute it under environmental disturbances, and strictly adhere to COLREGs. Crucially, it must also *adapt* the plan when new evidence appears without undermining the operator's understanding of the vehicle's intent. Reviews of USV mission planning consistently highlight that the most difficult challenges emerge not within individual algorithms, but at these interfaces between decision logic, safety, and remote operations (Xing et al., 2023) under harsh weather conditions.

While many studies treat components such as coverage planning and robust control in isolation, integrating them towards underwater object detection and identification type missions remains an open problem. Coverage and route planning methods have matured significantly and many studies combine global planning with local adjustment to handle obstacles and uncertainty (Teng et al., 2023). In parallel, Classical guidance and control methods such as line-of-sight tracking and model predictive control are widely used for robust path following under disturbances, often paired with separate collision-avoidance layers

(Perera et al., 2015). Although these methods are analyzable and often gives stability guarantees, hey become harder to apply in missions that must adapt to uncertain detections and changing priorities, and detailed discretization of dynamics and constraints quickly lead to high-dimensional models that are difficult to treat rigorously.

More recently, learning-based approaches, including deep reinforcement learning (DRL), have been proposed for collision avoidance and navigation under complex encounter dynamics, and encoding COLREGs in the reward design (Meyer et al., 2020). While promising, these methods face deployment barriers. Learned policies can be difficult to explain, and safety assurance remains challenging when decision logic is embedded in high-dimensional function approximators. Thus, such black-box decision process is often unacceptable in operational settings (Chen et al., 2021).

To bridge this gap, this paper proposes a layered mission-planning framework for underwater object detection and identification that prioritizes safety and auditability while enabling data-driven adaptability. Thus, as shown in Figure 1, a hybrid approach is taken: established low-level control and collision-avoidance modules provide a safe, analyzable backbone, while learning methods are used in a restricted advisory role to improve the timing and location of adaptive edits to the survey plan. Mission logic is expressed through an explicit Executive to preserve explainability and auditability.

The remainder of this paper is organized as follows. Section 2 reviews recent developments in mission-planning for autonomous marine systems, with emphasis on how established methods in decision logic, planning, and control, and safety enforcement are combined in practice. Section 3 then presents the proposed reference architecture in detail, describing the roles and interfaces of the executive, global planning, advisory learning, local tracking, and COLREG-aware risk and collision-avoidance layers. Section 4 outlines the current implementation direction on a ROC-supervised USV platform and the associated data products and workflows that support future surface and underwater vehicles in various sea trials.

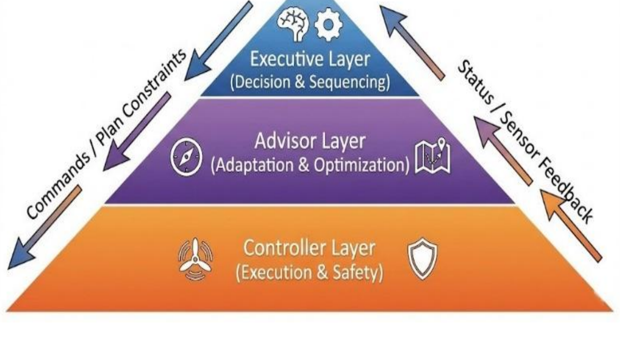

Figure 1. Mission Planning Framework - Layered Hierarchy

2 RELATED WORKS

The practical difficulty for mission planning for marine surface and underwater vehicles lies in composing multiple layers into a coherent system that can be deployed and supervised (Thompson and Guihen, 2019). Field demonstrations in underwater object detection style scenarios show that a mission can be degraded by intermittent sensing or loss of tracking, requiring explicit recovery behaviors and coordination procedures rather than only a better controller or planner (Ferri and Djapic, 2013). These observations motivate mission-planning frameworks that make the flow of intent and responsibility explicit, and that treat adaptation as an operational necessity rather than an optional feature.

To sequence phases and execution, mission intent must be represented in a form that supports both automation and supervision. One approach uses automated planning, where missions are encoded as symbolic constraints. For example, (Jang et al., 2022) has applied PDDL-based formalisms to underwater survey missions to handle complex resource constraints. While powerful, these symbolic approaches introduce significant modeling overhead and can obscure real-time system state from operators.

Alternative engineering-oriented approaches prioritize implementability and monitoring. "Rudimentary" mission-planning systems for autonomous surface ships illustrate how mission-level logic can be connected to an existing guidance and control stack that already contains path planning, path following, and collision avoidance (Hinostroza and Lekkas, 2022). Within this space, Behavior Trees (BTs) have become a standard representation due to their modularity and readability. BTs facilitate troubleshooting through explicit node statuses and provide a structured skeleton that can support limited learning without surrendering the control of mission sequencing (Iovino et al., 2022). Furthermore, large language models based techniques are also explored but there are still concerns regarding practical deployment and safety (Din et al., 2025).

Once mission logic is in place, seabed surveying in object detection and identification adds a specific planning demand: turning a survey polygon into an efficient coverage trajectory and updating it when new information appears. USV path and route planning surveys organize global planners into graph-search, sampling-based, optimization-based, and bio-inspired families, and many methods combine global planning with local adjustment to handle obstacles and uncertainty (Hashali et al., 2024). Long missions in underwear object detection and identification often require updates driven by detections, uncertainty, or operational constraints, and full replanning can disrupt mission continuity, confuse remote supervision, or impose computational cost at the wrong time. At the same time, purely local adjustments can preserve safety but erode coverage efficiency or produce behaviors that appear inconsistent with operator intent.

Once a route is produced, the mission architecture must still ensure tracking performance and coordination, especially when the mission involves repeated survey lines, formation behaviors, or constrained maneuvering. At the local level, USVs typically employ guidance and control strategies such as line-of-sight (LOS) guidance, PID control, or model predictive control (MPC) that track the planned trajectory while rejecting disturbances from waves and currents using state estimation techniques such as Kalman filter (Liu et al., 2022, Ahmed et al., 2024, Wang et al., 2024). Collision-avoidance modules are sometimes integrated as separate layers using velocity obstacles, model-predictive safety filters or rule-based COLREGs handling (Liu et al., 2023, Lu, 2025). While these components can often be analyzed individually, proving properties for the overall discrete-time closed loop becomes increasingly difficult as models are refined and interconnected. Furthermore, RL and fuzzy-logic based policies can generate complex avoidance maneuvers in simulation (Perera et al., 2012), however they struggle with interpretability and determining legality in multi-vehicle scenarios (Namazi and Perera, 2025).

Safe execution requires strict adherence to COLREGs during interactions. Research in this domain ranges from rule-based logic to DRL. While DRL policies can generate complex avoidance maneuvers in simulation, they struggle with interpretability and determining legality in multi-vehicle scenarios (Xu et al., 2022). Approaches that combine avoidance with fuzzy risk evaluation highlight that avoidance is fundamentally a risk decision (Perera and Soares, 2015); thus, the reasoning must be exposed rather than hidden inside opaque heuristics (Chen et al., 2021).

To mitigate this, recent research advocates for learning to augment rather than replace the executive, confining data-driven policies to bounded roles where they can be audited. Studies on regulatory readiness emphasize the need for clear alarm governance and explicit intervention rules, while risk assessment frameworks argue that mode switching must be structured and justified, particularly under communication limits (Fan et al., 2021). Furthermore, systematic reviews favor hybrid risk pipelines, using the outputs of qualitative tool for identifying unsafe control actions interaction analysis (e.g., STAMP) with probabilistic models such as Bayesian networks (Yuzui and Kaneko, 2025). Collectively, these works support architectures that keep decision authority legible and restrict learning to advisory roles, rather than placing opaque policies at the core of safety-critical control.

3 PROPOSED ARCHITECTURE

Figure 2 illustrates the proposed architecture for ROC-supervised seabed surveying on underwater object detection and identification with a USV and downstream task handoff to cooperating AUVs. The design is intentionally layered: it separates mission intent and sequencing from geometric planning, separates planning from execution, and keeps safety enforcement authoritative below any adaptive component. The framework targets coastal and harbor survey missions in which a ROC supervises a USV conducting wide-area MBES scanning, while a small team (3 to 4) of AUVs performs follow-up inspection. We assume polygon-defined survey areas at practical field scale, moderate currents and waves that require robust tracking, and communications that may degrade intermittently. The ROC–USV link supports periodic telemetry and command updates, whereas USV–AUV communication is treated as low-bandwidth and opportunistic, motivating intent-level tasking and compact reporting. The present scope considers a single USV coordinating multiple AUVs through task allocation and reporting protocols, rather than addressing full multi-AUV consensus and formation control.

The mission begins at the ROC, which is capable of directly communicating with USV. The ROC developed can perform multiple functions, supervise in navigation and send commands to the USV (Adnan et al., 2024a). Thus, the ROC can start the mission by providing a survey polygon and mission constraints such as desired line spacing, speed limits, expected endurance margins, and reporting preferences. At the top of the onboard stack sits the Mission Executive, which is responsible for sequencing mission phases and reacting to events. We implement this layer using a Behavior Tree structure that decomposes the mission into distinct phases, including system startup, transit, survey execution, adaptive inspection, and return to base. This ensures that the sequencing of the mission remains fully auditable and can be logged with reason codes.

BT root node represents the overall mission, with child subtrees corresponding to major phases: *System start-up and health checks*, *Survey polygon execution*, *Adaptive revisits* and *Return-to-base (RTB)*. Each leaf node encapsulates either a condition (e.g. “all mandatory sensors online”, “battery levels above threshold”, “communication link healthy”) or an action (e.g. “request polygon from ROC”, “initiate coverage plan”, “command RTB”). The BT structure enforces a clear sequencing and recovery strategy: failures in sensor checks trigger fallback nodes that request operator intervention, while run-time events such as high-risk scores or low battery trigger transitions from the survey phase to adaptive revisits or RTB.

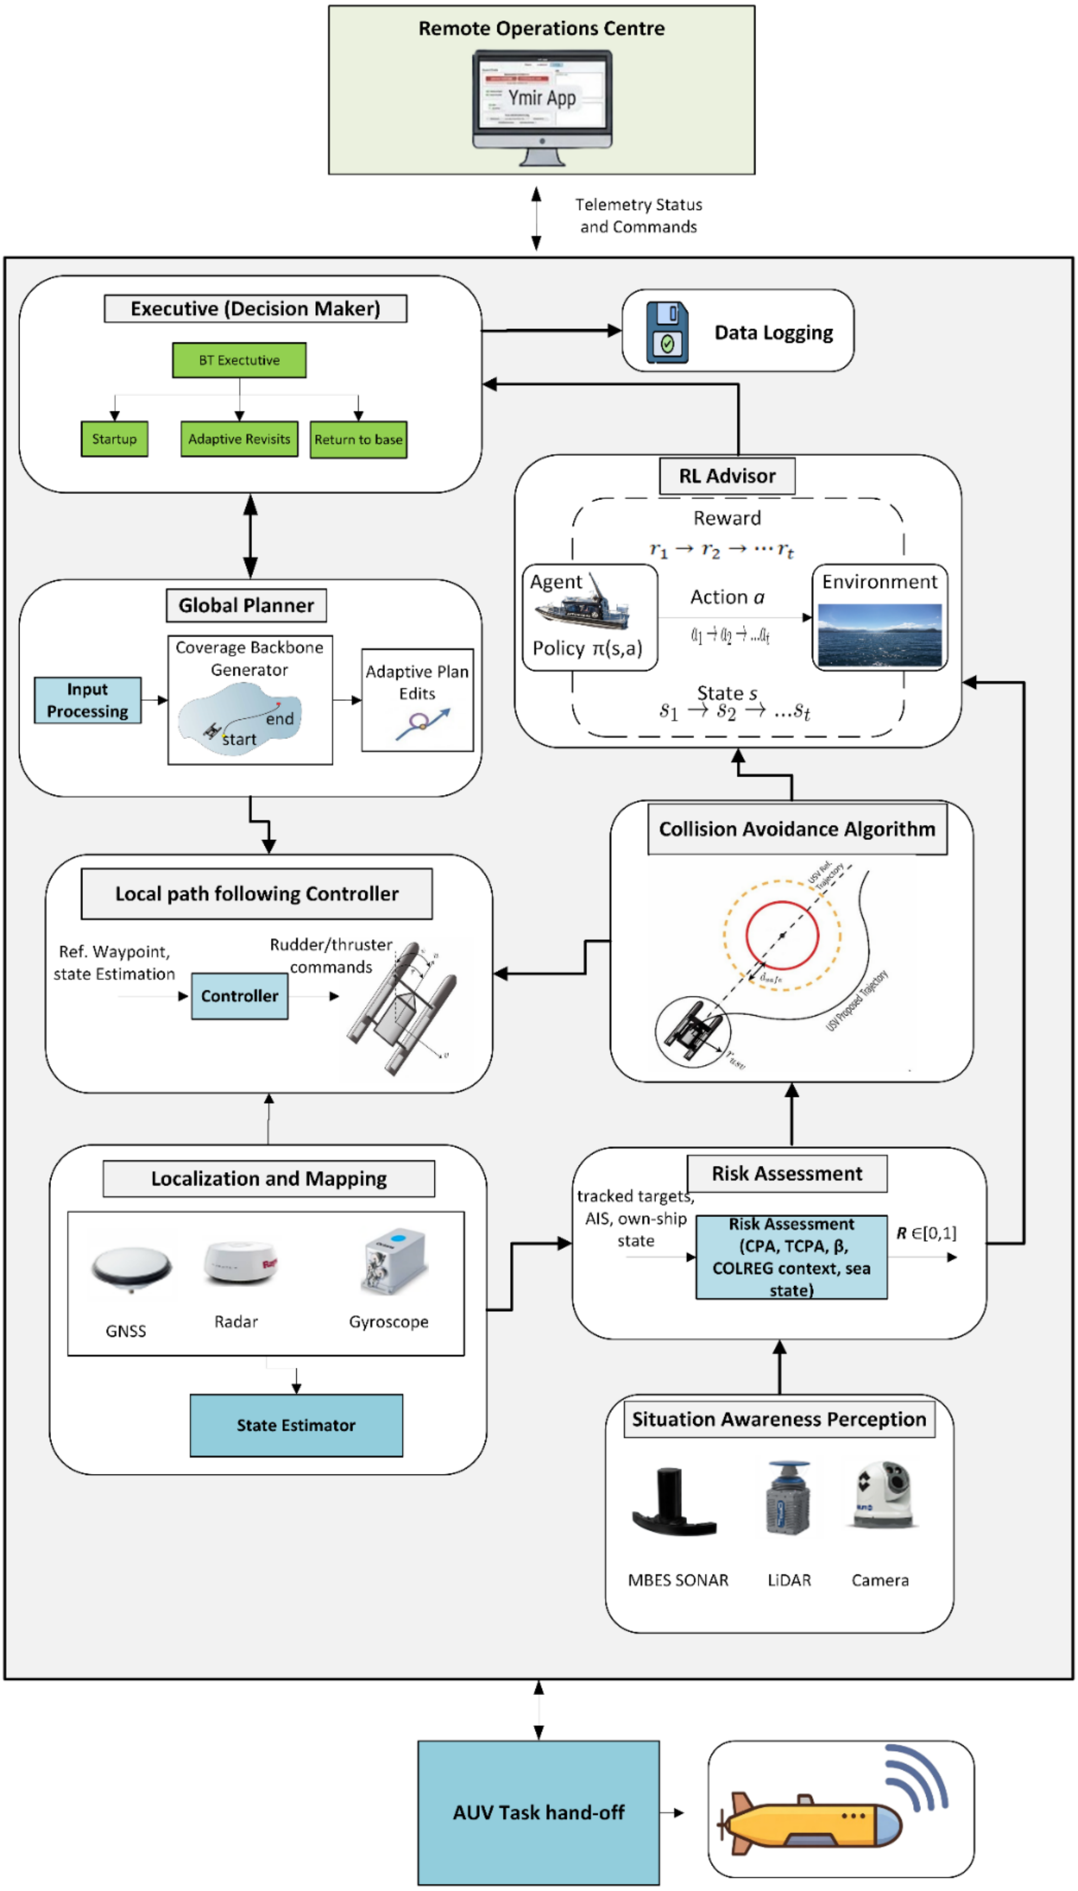

Figure 2. Detailed Architecture for Mission Planning of USV

Beneath the Executive, the Global Planner maintains the geometric definition of the mission. Coverage path planning for polygonal regions is commonly realized through lawnmower patterns derived from the mission polygon and sensor footprint, often after a decomposition and line-ordering step, with subsequent optimization for cost, time, or environmental effects (Hashali et al., 2024). These methods provide an efficient and predictable baseline for hydrographic and bathymetric surveys, but they typically treat mission-level updates as either a re-run of the planner with updated constraints or as incremental additions that are not explicitly constrained from an operator-supervision perspective. Rather, this layer employs a backbone-and-edit approach. The planner generates an initial coverage pattern based on the operator-defined polygon and sensor constraints. During execution, this stable backbone is preserved while dynamic adaptations are introduced as bounded edits. When a region requires higher resolution data or a revisit, the planner splices a temporary sub-path into the lane sequence.

Conceptually, the planner maintains an ordered list of path segments $\{\gamma_i\}$. When triggered by an event (e.g. a high-confidence detection cluster or a region with low coverage density), a local edit operator replaces a subsequence $\gamma_{k:k+m}$ with a modified subsequence $\tilde{\gamma}_{k:k+\tilde{m}}$ that increases information gain while keeping the rest of the plan intact. Each edit is constrained to a local window and must rejoin the backbone within a short horizon, which preserves mission continuity for ROC supervision and reduces disruption to the executive state. Smooth transition segments connect edits to the main path so the vehicle can deviate for inspection and rejoin the primary survey route without discarding the global plan.

Adaptive edits are triggered and mediated by the Executive, supported by an advisory learning module. The advisor is designed to suggest when and where adaptations are likely to be beneficial, while remaining outside the safety-critical control loop. This constraint is deliberate. Edits are triggered by explicit operational cues, such as clustered high-confidence detections from the onboard survey pipeline, coverage-density deficits, updated constraints from the risk layer, communication-mode changes, or operator requests. An edit is accepted only after a feasibility-and-safety gate that checks time and energy budgets, curvature and controller limits, and predicted collision or rule-violation risk (e.g., via CPA/TCPA thresholds and COLREG-constrained avoidance). Reinforcement learning can be applied to maritime navigation and collision avoidance, including attempts to encode COLREG considerations into the learning objective, and these studies demonstrate that learned policies can generate context-dependent behaviors in simulated encounter settings.

The advisory module observes a compact mission state vector s_t that includes features such as current coverage density around the vessel, detection confidence and clustering statistics for potential seabed objects, estimated risk levels, any potential obstacles, remaining energy and time, and environmental conditions. Based on this state, the advisor outputs an action a_t from a small set, such as continuing the baseline plan, creating a rescan box around hotspot i, densifying a lane segment, or proposing an alternate local route around an obstruction. The Executive can accept, defer, or reject these recommendations based on explicit rules and mission constraints. In effect, learning is used to improve timing and prioritization of edits, while the deterministic backbone remains responsible for legality and safety.

Once a path segment is authorized, it is passed to the Local Tracking Layer for execution. It translates the current global segment into a reference trajectory and generates actuator commands under disturbances. This layer utilizes a Model Predictive Control framework to generate thrust and rudder commands that minimize cross-track error while respecting the vessel's dynamic limitations. The use of predictive control allows the system to handle environmental disturbances such as currents and wind drift explicitly. The controller gets input from the State Estimator module, which uses techniques like Kalman filtering and Moving Horizon Estimation techniques to estimate the true states of the USV using sensors (GNSS, IMU, Gyroscope) output. The tracker operates as a dependable service that attempts to follow the commanded reference trajectory regardless of the higher-level mission context, ensuring a clean separation of concerns between planning and control.

Safety is enforced by a dedicated risk assessment and collision-avoidance layer that monitors the maritime environment and constrains execution to remain COLREG compliant. This layer consumes tracked targets, own-ship state, and perception outputs, and it computes risk indicators that can be logged and reported to the ROC. Risk can be assessed using structured heuristics or fuzzy evaluation mechanisms that combine encounter geometry and time-to-closest-approach style variables into interpretable risk levels, which can then trigger avoidance behavior in a traceable manner. In the proposed architecture, the collision-avoidance layer has overriding authority over the tracker when necessary. Any override is reported upward to the Executive as a structured event so that the mission state remains coherent. After an avoidance maneuver, the system performs a controlled rejoin to the global plan, maintaining continuity of the survey rather than leaving the Executive in an ambiguous state about what segment is currently being executed.

To support follow-up identification, the USV also serves as the coordination and communications bridge between the ROC and one or more AUVs. Candidate detections from the MBES are packaged onboard into compact inspection tasks, each containing a region of interest, priority, uncertainty, and basic execution constraints, and are transmitted using intent-level messages suited to low-bandwidth underwater links. In case multiple AUVs are available, well-established task-allocation methods and multi-agent consensus algorithms can be used to assign each inspection task to a single vehicle based on simple availability and cost indicators returned by the AUVs (Wang et al., 2022). The AUV returns acknowledgements, health and status, and a compact inspection summary. The USV then aggregates acknowledgments and inspection summaries for relay back to the ROC.

The architecture is designed to remain functional even under the intermittent communication conditions typical of maritime environments by adapting its reporting to link quality. During normal operation, the vessel streams comprehensive telemetry and sensor snippets to ROC. In degraded conditions, the system switches to intent-based reporting, transmitting only high-level status updates and event logs while buffering heavy sensor data onboard. This ensures that the remote operator maintains an accurate mental model of the vehicle's phase and intent, even when real-time monitoring is not possible. All detections are cataloged onboard with associated confidence scores and metadata, ready for transmission to cooperative underwater assets or post-mission analysis.

During harbor and sea trials, the framework will be instrumented so that each layer produces synchronized, time-stamped logs that support quantitative evaluation of mission outcomes and operator supervision. The trials will quantify survey effectiveness over the operator-defined polygon using footprint-based coverage metrics (Lu et al., 2024). Coverage completeness will be reported as $C = A_{\text{covered}}/A_{\text{poly}}$, where A_{covered} is the polygon area receiving MBES footprint at the required line spacing and A_{poly} is the mission polygon area. Survey redundancy will be reported through overlap or re-coverage fraction, since MBES survey practice commonly trades overlap against time and quality. The cost of adaptation will be measured by edit latency t_{edit}, defined as the time to generate, validate, and accept a bounded edit, and by relative path overhead $\Delta L = (L - L_0)/L_0$, where $L_0$is the backbone length and $L$is the executed length after edits.

Safety performance will be summarized from encounter logs using closest-point-of-approach measures, including minimum separation $d_{\min}$, DCPA/TCPA where available, and counts of CPA-threshold events that triggered avoidance or safety overrides (Qu and Cai, 2022). The operational impact of safety actions will be reported as the frequency and

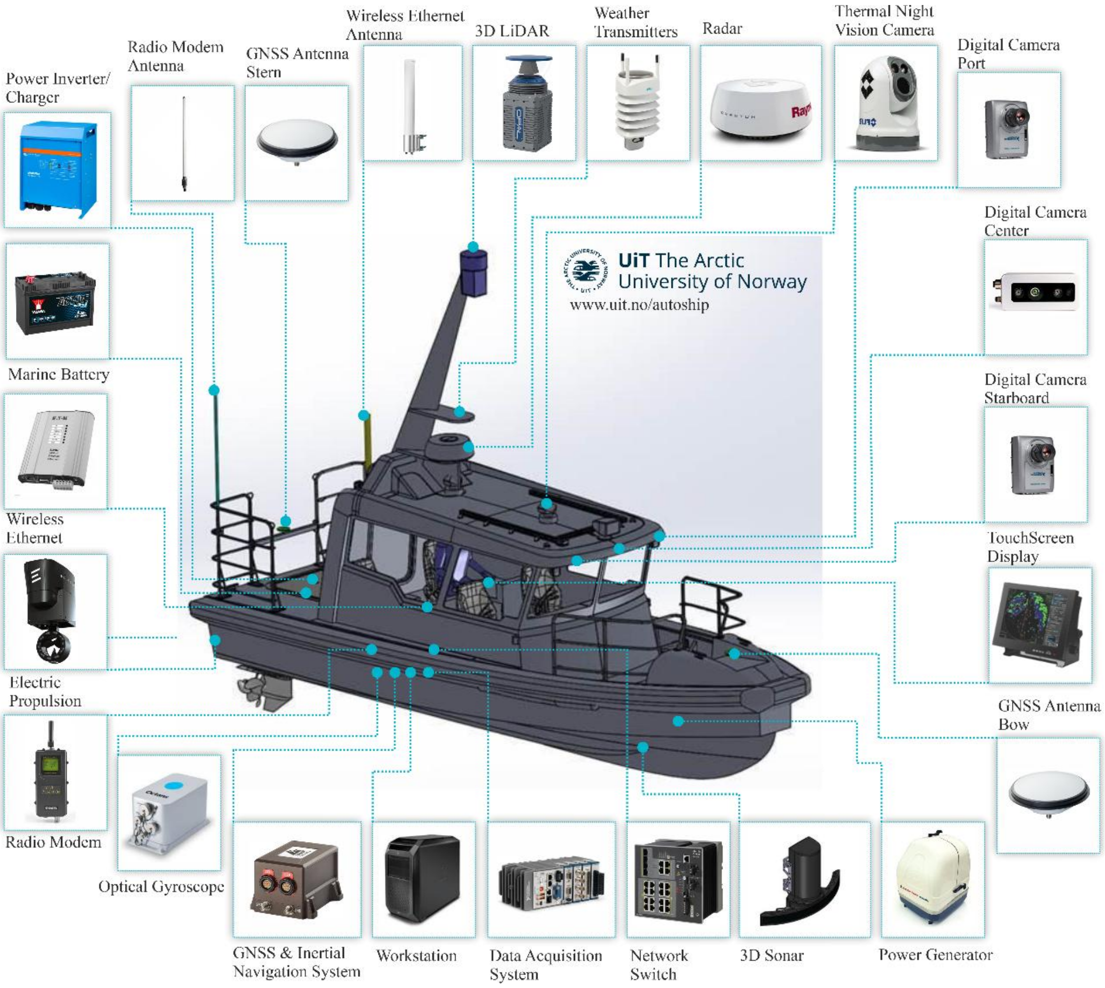

Figure 3. Installed sensors on Ymir surface vessel

duration of override episodes and the time-to-rejoin the path after avoidance. ROC-supervised operation under intermittent links will be evaluated using reporting-mode statistics and timeliness, including event-report latency, time spent in each reporting mode, and buffered-data backlog during degraded communication. If AUV inspection tasks are issued, the trials will report task acknowledgement rate, time-to-inspect a region of interest, and completeness of the returned inspection summary under constrained underwater communications, using standard reliability and delay measures as supporting indicators.

4 REMOTE OPERATION CENTRE AND USV COMMUNICATION PLATFORM

The proposed framework is being implemented on the UiT research vessel Ymir and an associated Remote Operation Centre (Adnan and Perera, 2024). Ymir is an 8 meter vessel with hybrid-electric platform to support extended autonomous operations with lower vibration and precise low-speed control. The propulsion system consists of a marine battery bank driving an electric propulsion motor, supported by an onboard power generator and inverter/charger system for range extension. The vessel is equipped with a comprehensive sensor suite designed to feed the perception and safety layers described in Section 3. The complete sensor modules are illustrated in Figure 3.

Navigation relies on a high-precision GNSS/INS solution with dual antennas (bow and stern) and an optical gyroscope to ensure heading accuracy even during dynamic maneuvers. For surface perception and collision avoidance, the vessel integrates a Raymarine radar, a 3D LiDAR, and a thermal night-vision camera, providing redundant detection modalities for obstacle tracking. The MBES is deployed for bathymetry and real-time underwater object detection and identification. All sensor data is aggregated by an onboard high-performance workstation and a dedicated data acquisition system, which will host the Mission Executive, the RL Advisor, and the local controllers.

To support remote supervision, a customized communication architecture links the vessel to the ROC via a combination of high-bandwidth wireless Ethernet (for near-shore operations) and a robust long-range radio modem. At the ROC, a graphical user interface enables operators to monitor vessel position, visualize sensor status and streams, and control which sensors are active for a given mission. The GUI includes dedicated panels for starting and stopping sensors and logging sessions, data collection, recording and playback. The GUI can display the data from different sensors installed in the USV, such as digital and night-vision cameras, sonar and lidar views, and other navigation data; start/stop the sensors and record the data, as shown in Figure 4.

This interface can be further extended to display other relevant modules and sensors as per the mission planning requirement. GUI also displays the live path route of the USV, and the latest navigation data as shown in Figure 5, such that in case of communication loss with the USV, the latest position of USV is available at the OCC. The ROC thus provides both a human-in-the-loop supervision layer and a data source for training and validating advisory learning components. Furthermore, the ROC has the capability to directly control the different sensors and modules installed in the USV, and complete log of the USV status is displayed as well as recorded, as shown in Figure 6.

5 CONCLUSION

This paper proposed a hierarchical mission-planning framework for ROC-supervised seabed surveys, where a USV performs wide-area sensing and coordination and AUVs support follow-up inspection. The framework treats mission planning as a systems problem by combining an explicit executive for mission phases and recovery behaviors, a coverage planner

Figure 4. Data Collection System on ROC

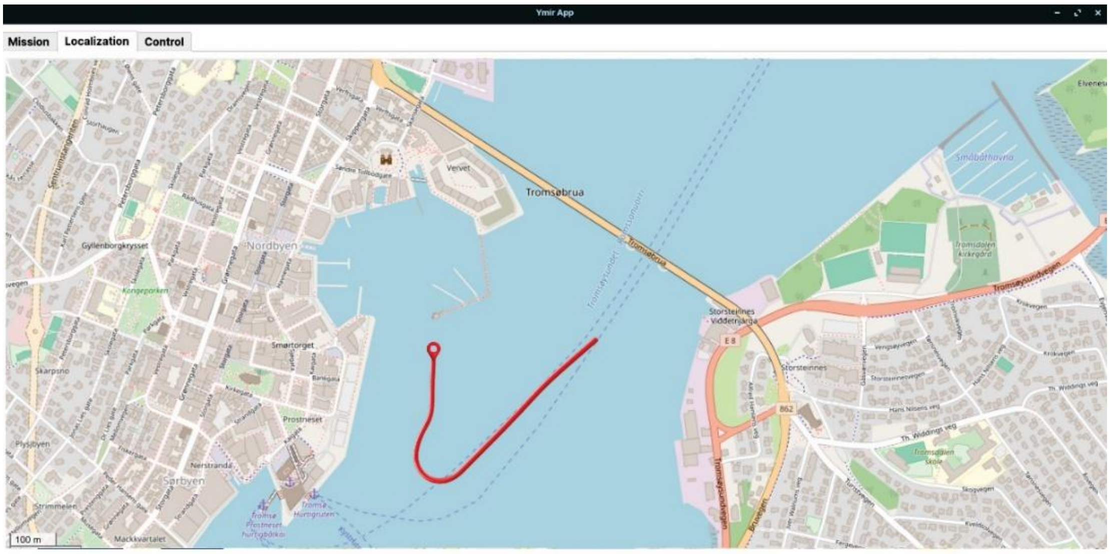

Figure 5. USV localization and visualization on ROC

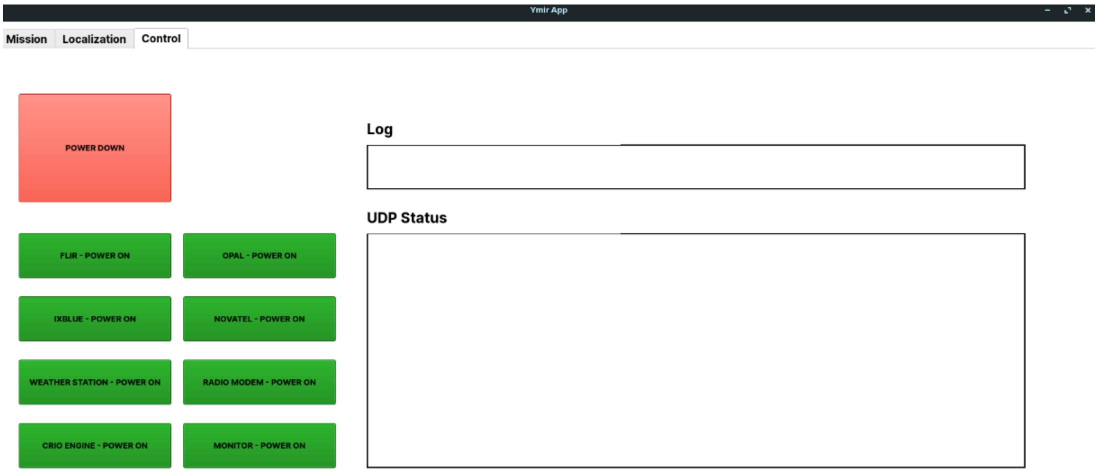

Figure 6. Mission Control Module of ROC

that supports bounded local edits instead of re-peated full replanning, and authoritative safety enforcement through risk assessment and collision avoidance. Learning is used only in an advisory role to recommend when and where edits are beneficial, while decisions and actuation remain auditable and logged. The paper also outlined implementation readiness through the evolving Ymir sensor stack and ROC workflow. Future work will complete the architecture and validate it in a selected sea trials in Tromsø, Norway, including degraded-communication operation and structured handoff of targets to underwater assets.

Acknowledgements: This research is a part of the MMinE-SwEEPER project and has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101167839.

REFERENCES

Adnan, M. & Perera, L. 2024. Navigational support framework for maritime autonomous surface ships under onshore operation centers. *Advances in Maritime Technology and Engineering.* CRC Press.

Adnan, M. & Perera, L. P. Virtual Reality Based Simulator Platform Utilization for Training Remote Vessel Operators. International Congress of the International Maritime Association of the Mediterranean, 2025. Springer, 350–361.

Adnan, M., Perera, L. P., Tusher, H. M. & Nazir, S. Functional requirements for onshore operation centers to support remotely operated ships. ISOPE International Ocean and Polar Engineering Conference, 2024a. ISOPE, ISOPE–I–24–605.

Adnan, M., Wang, Y. & Perera, L. P. C. 2024b. Role of Onshore Operation Centre and Operator in Remote Controlled Autonomous Vessels Operation.

Ahmed, S. H., Zhao, M. & Li, H. Cooperative Path Following Control of Unmanned Surface Vehicles Using Model Predictive Control. 2024 IEEE 7th International Conference on Industrial Cyber-Physical Systems (ICPS), 2024. IEEE, 1–6.

Chen, Y.-Y., Ellis-Tiew, M.-Z., Chen, W.-C. & Wang, C.-Z. 2021. Fuzzy risk evaluation and collision avoidance control of unmanned surface vessels. *Applied Sciences,* 11**,** 6338.

Din, M. U., Akram, W., Bakht, A. B., Dong, Y. & Hussain, I. Maritime Mission Planning for Unmanned Surface Vessel using Large Language Model. 2025 IEEE

International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR), 2025. IEEE, 1–6.

Fan, C., Montewka, J. & Zhang, D. 2021. Towards a framework of operational-risk assessment for a maritime autonomous surface ship. *Energies,* 14**,** 3879.

Ferri, G. & Djapic, V. Adaptive mission planning for cooperative autonomous maritime vehicles. 2013 IEEE International Conference on Robotics and Automation, 2013. IEEE, 5586–5592.

Hashali, S. D., Yang, S. & Xiang, X. 2024. Route planning algorithms for unmanned surface vehicles (USVs): a comprehensive analysis. *Journal of Marine Science and Engineering,* 12**,** 382.

Hinostroza, M. A. & Lekkas, A. M. 2022. A rudimentary mission planning system for marine autonomous surface ships. *IFAC-PapersOnLine,* 55**,** 196–203.

Iovino, M., Scukins, E., Styrud, J., Ögren, P. & Smith, C. 2022. A survey of behavior trees in robotics and ai. *Robotics and Autonomous Systems,* 154**,** 104096.

Jang, J., Do, H. & Kim, J. 2022. Mission planning for underwater survey with autonomous marine vehicles. *Journal of Ocean Engineering and Technology,* 36**,** 41–49.

Liu, W., Qiu, K., Yang, X., Wang, R., Xiang, Z., Wang, Y. & Xu, W. 2023. COLREGS-based collision avoidance algorithm for unmanned surface vehicles using modified artificial potential fields. *Physical Communication,* 57**,** 101980.

Liu, Z., Song, S., Yuan, S., Ma, Y. & Yao, Z. 2022. ALOS-based USV path-following control with obstacle avoidance strategy. *Journal of Marine Science and Engineering,* 10**,** 1203.

Lu, J. 2025. Research on Intelligent Collision Avoidance Decision-Making Algorithm for Multi-Unmanned Surface Vehicles (USVs) Based on COLREGs (March 2025). *IEEE Access*.

Lu, Y., Xu, J., Zhong, Y. & Lin, H. 2024. Enhancing Marine Topography Mapping: A Geometrically Optimized Algorithm for Multibeam Echosounder Survey Efficiency and Accuracy. *Applied Sciences,* 14**,** 8875.

Meyer, E., Heiberg, A., Rasheed, A. & San, O. 2020. COLREG-compliant collision avoidance for unmanned surface vehicle using deep reinforcement learning. *Ieee Access,* 8**,** 165344–165364.

Namazi, H. & Perera, L. P. Adaptive Neuro-fuzzy Inference for Predicting Control Inputs in Maritime Autonomous Surface Ships Maneuvers. International Congress of the International Maritime Association of the Mediterranean, 2025. Springer, 221–234.

Perera, L. P., Carvalho, J. & Soares, C. G. 2012. Intelligent ocean navigation and fuzzy-Bayesian decision/action formulation. *IEEE journal of Oceanic Engineering,* 37**,** 204–219.

Perera, L. P., Ferrari, V., Santos, F. P., Hinostroza, M. A. & Soares, C. G. 2015. Experimental evaluations on ship autonomous navigation and collision avoidance by intelligent guidance. *IEEE Journal of Oceanic Engineering,* 40**,** 374–387.

Perera, L. P. & Soares, C. G. 2015. Collision risk detection and quantification in ship navigation with integrated bridge systems. *Ocean Engineering,* 109**,** 344–354.

Qu, Y. & Cai, L. 2022. Real-time emergency collision avoidance for unmanned surface vehicles with COLREGS flexibly obeyed. *Journal of Marine Science and Engineering,* 10**,** 2025.

Teng, S., Hu, X., Deng, P., Li, B., Li, Y., Ai, Y., Yang, D., Li, L., Xuanyuan, Z. & Zhu, F. 2023. Motion planning for autonomous driving: The state of the art and future perspectives. *IEEE Transactions on Intelligent Vehicles,* 8**,** 3692–3711.

Thompson, F. & Guihen, D. 2019. Review of mission planning for autonomous marine vehicle fleets. *Journal of Field Robotics,* 36**,** 333–354.

Wang, S., Liu, Y., Qiu, Y. & Zhou, J. 2022. Consensus-based decentralized task allocation for multi-agent systems and simultaneous multi-agent tasks. *IEEE Robotics and Automation Letters,* 7**,** 12593–12600.

Wang, Y., Perera, L. P. & Batalden, B.-M. Adaptive Kalman Filter-based Estimator with Sea Trail Data to Calculate Ship States in Complex Navigation Conditions. ISOPE International Ocean and Polar Engineering Conference, 2024. ISOPE, ISOPE–I–24–536.

Xing, B., Yu, M., Liu, Z., Tan, Y., Sun, Y. & Li, B. 2023. A review of path planning for unmanned surface vehicles. *Journal of Marine Science and Engineering,* 11**,** 1556.

Xu, X., Cai, P., Ahmed, Z., Yellapu, V. S. & Zhang, W. 2022. Path planning and dynamic collision avoidance algorithm under COLREGs via deep reinforcement learning. *Neurocomputing,* 468**,** 181–197.

Yuzui, T. & Kaneko, F. 2025. Toward a hybrid approach for the risk analysis of maritime autonomous surface ships: a systematic review. *Journal of Marine Science and Technology,* 30**,** 153–176.

Zhao, L., Bai, Y. & Paik, J. K. 2024. Optimal coverage path planning for USV-assisted coastal bathymetric survey: Models, solutions, and lake trials. *Ocean Engineering,* 296**,** 116921.